BURROUGHS WELLCOME FUND

DeepMind

GATSBY

SIMONS FOUNDATION

Program Summary

All times are EST (Eastern Standard Time, GMT-5)

Tuesday, 23 February

11:00a	**Cosyne Tutorial: Recurrent Neural Networks for Neuroscience**
03:00p	**Paths for leadership in promoting the careers of scientists**

Wednesday, 24 February

08:00a	**Poster Session 1a**
10:00a	**Session 1:** 3 accepted talks
11:40a	**Session 2:** 4 accepted talks
01:00p	**Simons Foundation Bridge to Independence Awards: Facilitating the transition from postdoc to junior faculty**
02:00p	**Session 3:** 3 accepted talks
03:00p	**Poster Session 1b**
05:00p	**Mini-symposium on the history of race and racism in science**
07:00p	**Poster Session 1c**

Thursday, 25 February

08:00a	**Poster Session 2a**
10:00a	**Session 4:** 3 accepted talks
11:40a	**Session 5:** 4 accepted talks
02:00p	**Session 6:** 3 accepted talks
03:00p	**Poster Session 2b**
05:00p	**Expanding horizons: Diverse perspectives in computational neuroscience**
07:00p	**Poster Session 2c**

Friday, 26 February

08:00a	**Poster Session 3a**
10:00a	**Session 7:** 3 accepted talks
11:40a	**Session 8:** 4 accepted talks
02:00p	**Session 9:** 3 accepted talks
03:00p	**Poster Session 3b**
07:00p	**Poster Session 3c**

About Cosyne

The annual Cosyne meeting provides an inclusive forum for the exchange of experimental and theoretical/computational approaches to problems in systems neuroscience.

To encourage interdisciplinary interactions, the main meeting is arranged in a single track. A set of invited talks are selected by the Executive Committee and Organizing Committee, and additional talks and posters are selected by the Program Committee, based on submitted abstracts.

Cosyne topics include (but are not limited to): neural basis of behavior, sensory and motor systems, circuitry, learning, neural coding, natural scene statistics, dendritic computation, neural basis of persistent activity, nonlinear receptive field mapping, representations of time and sequence, reward systems, decision-making, synaptic plasticity, map formation and plasticity, population coding, attention, machine learning for neuroscience, and computation with spiking networks. Participants include pure experimentalists, pure theorists, and everything in between.

Cosyne 2021 Leadership

Organizing Committee

General and Program Chairs

Anne-Marie Oswald (University of Pittsburgh) and Srdjan Ostojic (Ecole Normale Superieure Paris)

Publicity Chair

Adam Calhoun (Princeton University)

Diversity Chairs

Eva Dyer (Georgia Tech, Emory University) and Eric Shea-Brown (University of Washington)

Tutorial Chair

Il Memming Park (Stony Brook University)

Development Chair

Michael Long (New York University)

Media Chair

Carlos Stein Brito (EPFL)

Executive Committee

Stephanie Palmer (University of Chicago)

Zachary Mainen (Champalimaud)

Alexandre Pouget (University of Geneva)

Anthony Zador (Cold Spring Harbor Laboratory)

Program Committee

Anne-Marie Oswald (University of Pittsburgh), co-chair
Srdjan Ostojic (Ecole Normale Superieure Paris), co-chair
Athena Akrami (University College London)
Demba Ba (Harvard University)
Omri Barak (Technion)
Brice Bathellier (Paris)
Yoram Burak (Hebrew University)
Christine Constantinople (New York University)
Saskia De Vries (Allen Institute for Brain Science)
Victor de Lafuente (National Autonomous University of Mexico)
Jan Drugowitsch (Harvard University)
Tatiana Engel (Cold Spring Harbor Laboratory)
Annegret Falkner (Princeton University)
Kevin Franks (Duke University)
Tim Hanks (University of California, Davis)
Ann Hermundstad (Janelia Farm Research Campus)
Andrew Leifer (Princeton University)
Sukbin Lim (New York University Shanghai)
Scott Linderman (Stanford University)
Mackenzie Mathis (Ecole polytechnique federale de Lausanne)
Ida Momennejad (Microsoft)
Cris Niell (University of Oregon)
Il Memming Park (Stony Brook)
Adrien Peyrache (McGill University)
Kanaka Rajan (Icahn School of Medicine at Mount Sinai)
Robert Rosenbaum (University of Notre Dame)
Christina Savin (New York University)
Marshall Shuler (Johns Hopkins University)
Daniela Vallentin (Max Planck Institute, Munich)

Cosyne 2021 reviewers

Alireza Alemi, Cristiano Alessandro, Andy Alexander, Mikio Aoi, Diego Armando Pacheco Pinedo, Florencia Assaneo, Edoardo Balzani, Joao Barbosa, Tania Barkat, Jérémie Barral, Andrea Barreiro, Helen Barron, Eleanor Batty, Jeff Beck, Mor Ben-Tov, Kevin Berlemont, Gordon Berman, Johannes Bill, Kevin Bolding, Kathryn Bonnen, Flora Bouchacourt, Braden Brinkman, Iva Brunec, Denise Cai, Adam Calhoun, Dario Campagner, Ioana Carcea, Luis Carrillo-Reid, Sourish Chakravarty, Chandramouli Chandrasekaran, Adam Charles, Rishidev Chaudhuri, Xiaomo Chen, SueYeon Chung, Ruben Coen-Cagli, Julia Cox, Matt Creamer, Ran Darshan, Ian Davison, Daniel Denman, Emily Dennis, Erin Diel, Long Ding, Mario Dipoppa, May Dobosiewicz, Chunyu A Duan, Lea Duncker, Rodrigo Echeveste, Seth Egger, Valerie Ego-Stengel, Tahra Eissa, Ahmed El Hady, Sami El-Boustani, Ben Engelhard, Rainer Engelken, Jeffrey Erlich, Katie Ferguson, Andrew Fink, Walter Fischler, Alexander Fleischmann, Francisco Flores, Thomas Frank, Tomoki Fukai, Izumi Fukunaga, Gabrielle Girardeau, Joshua Glaser, Yi Gu, Grigori Guitchounts, Gabrielle Gutierrez, Ranier Gutierrez, Hannah Haberkern, Ralf Haefner, Steffen Hage, Arif Hamid, Gregory Handy, Guillaume Hennequin, Yu Hu, Chengcheng Huang, Ramakrishnan Iyer, Simon Jacob, Elina Jacobs, Xiaoxuan Jia, Kresimir Josic, Jonathan Kadmon, Gary Kane, Nikolaos Karalis, Matthias Kaschube, Loukia Katsouri, Alfred Kaye, Alex Kell, Ann Kennedy, Sepiedeh Keshavarzi, Zachary Kilpatrick, Christopher Kim, Sung Soo Kim, Sue Ann Koay, Mihaly Kollo, Charles Kopec, Kishore Kuchibhotla, Guillaume Lajoie, Kaushik Lakshminarasimhan, Angela Langdon, Kenneth Latimer, Yaara Lefler, Luis Lemus, Anna Levina, Songting Li, Laureline Logiaco, Matthew Lovett-Barron, Emilie Mace, Noa Malem-Shinitski, Christian Marton, Paul Masset, Francesca Mastrogiuseppe, Luca Mazzucato, Samuel McDougle, Matthew McGinley, Jenna McHenry, Jorge Mejias, Raoul-Martin Memmesheimer, Hugo Merchant, Leenoy Meshulam, Florent Meyniel, Daniel Millman, Aaron Milstein, Wiktor Mlynarski, Joseph Monaco, Ruben Moreno Bote, Malavika Murugan, Samuel Muscinelli, Vijay Mohan K Namboodiri, Devika Narain, Marcella Noorman, Cian O'Donnell, Michael Okun, Freyja Olafsdottir, Azahara Oliva, Amy Orsborn, Hysell Oviedo, David Owald, Nancy Padilla Coreano, Marino Pagan, Agostina Palmigiano, Cengiz Pehlevan, Matthew Perich, Catherine Perrodin, Megan Peters, Warren Pettine, Alex Piet, Lucas Pinto, Joshua Plotkin, Maximilian Puelma Touzel, Angela Radulescu, Srikanth Ramaswamy, Erin Rich, Marion Rouault, Lisa Roux, Alon Rubin, Mikail Rubinov, Peter Rudebeck, Pavel Rueda Orozco, Caroline Runyan, Shreya Saxena, Steffen Schneider, Carl Schoonover, Tilo Schwalger, Won Mok Shim, Hideaki Shimazaki, Joshua Siegle, Spencer Smith, Andrew Song, Kimberly Stachenfeld, Merav Stern, Weinan Sun, Vladislav Susoy, David Sussillo, Emily Sylwestrak, Sina Tafazoli, Dougal Tervo, Kay Thurley, Balazs Ujfalussy, Julia Veit, Vivek Venkatachalam, Siwei Wang, Michael Wehr, Xue-Xin Wei, Matthew Whiteway, Alex Williams, Marina Wosniack, Qiaojie Xiong, Guangyu Robert Yang, Manyi Yim, Fleur Zeldenrust, Yuan Zhao, Petr Znamenskiy

Special thanks to Titipat Achakulvisut, Daniel Acuna, and Konrad Kording for writing and managing the automated software for reviewer abstract assignment.

Conference Support

Administrative Support, Registration, Hotels

Leslie Weekes, Cosyne

Social media policy

Cosyne encourages the use of social media before, during, and after the conference, so long as it falls within the following rules:

- Do not capture or share details of any unpublished data presented at the meeting.
- If you are unsure whether data is unpublished, check with the presenter before sharing the information.
- Respect presenters' wishes if they indicate the information presented is not to be shared.

Stay up to date with Cosyne 2021 #cosyne21

Cosyne Code of Conduct

Purpose

At Cosyne, we strive for open and honest intellectual debate as part of a welcoming and inclusive atmosphere. This requires a community and an environment that recognizes and respects the inherent worth of every person.

Sources

This code of conduct is based on standards and language set at other meetings, whose organizing boards convened special working groups of scientific and legal experts to set their policies. We follow, in particular, those guidelines established for the Gordon Research Conferences, the Society for Neuroscience Annual Meeting, and NeurIPS.

The following code of conduct has been adapted from:

https://www.grc.org/about/grc-policies-and-legal-disclaimers

https://www.sfn.org/Membership/Professional-Conduct/Code-of-Conduct-at-SfN-Events

https://nips.cc/public/CodeOfConduct

Other online resources:

http://changingourcampus.org

https://www.sfn.org/Membership/Professional-Conduct/SfN-Ethics-Policy

Responsibilities

All participants, volunteers, organizers, reviewers, speakers, sponsors, and volunteers (referred to as "Participants" collectively throughout this document) at our Conference, workshops, and Conference-sponsored social events—are required to agree with this Code of Conduct both during an event and on official communication channels, including social media.

Sponsors are equally subject to this Code of Conduct. In particular, sponsors should not use images, activities, or other materials that are of a sexual, racial, or otherwise offensive nature. This code applies both to official sponsors as well as any organization that uses the Conference name as branding as part of its activities at or around the Conference.

Organizers will enforce this Code, and it is expected that all Participants will cooperate to help ensure a safe and inclusive environment for everyone.

Policy

The conference commits itself to providing an experience for all Participants that is free from the following:

Harassment, bullying, discrimination which includes but is not limited to:

- Offensive comments related to age, race, religion, creed, color, gender (including transgender/gender identity), sexual orientation, medical condition, physical or intellectual disability, pregnancy, or medical conditions, national origin or ancestry.
- Intimidation, personal attacks, harassment, unnecessary disruption of talks or other conference events.

Inappropriate or unprofessional behavior that interferes with another's full participation including:

- Sexual harassment, stalking, following, harassing photography or recording, inappropriate physical contact, unwelcome attention, public vulgar exchanges, derogatory name-calling, and diminutive characterizations.
- Use of images, activities, or other materials that are of a sexual, racial, or otherwise offensive nature that may create an inappropriate or toxic environment.
- Disorderly, boisterous, or disruptive conduct including fighting, coercion, theft, damage to property, or any mistreatment or non-businesslike behavior towards participants.
- "Zoom bombing" or any virtual activity that is not related to the topic of discussion which detracts from the topic or the purpose of the program. This includes inappropriate remarks in chat areas as deemed inappropriate by presenters/monitors/event leaders.

Scientific misconduct: including fabrication, falsification, or plagiarism of paper submissions or research presentations, including demos, exhibits or posters. Cosyne asks each session chair and organizing and reviewing committee member to promote rigorous analysis of all science presented for or at the meeting in a manner respectful to all attendees.

This Code of Conduct applies to the actual meeting sites and Conference venues where Cosyne business is being conducted, including physical venues, online venues, and official virtual engagement platforms, including video, virtual streaming, and chat-based interactions. Cosyne is not responsible for non-sponsored activity or behavior that may occur at non-sponsored locations such as hotels, restaurants, or physical or virtual locations not otherwise a sanctioned space for sponsored events. *Nonetheless, any issues brought to the Hotline Relations Counselors will be explored.* Moreover, Cosyne will not actively monitor social media platforms but will follow up on issues of harassment and violations of the code of conduct that occur on those platforms that are specifically related to the Cosyne program, during the course of Cosyne, if and when they are brought to our attention.

Complaint reporting

The Conference encourages all Participants to immediately report any incidents of discrimination,harassment, unprofessional conduct, and/or retaliation so that complaints can be quickly and fairly resolved. There will be no retaliation against any Participant who brings a complaint or submits an incident report in good faith or who honestly assists in investigating such a complaint. If you have concerns related to your participation/interaction at the Conference or Conference sanctioned events, or observe someone else's difficulties, or have any other concerns you wish to share, please write to *CosyneHotline@gmail.com* or by calling the *Cosyne Hotline phone number at +1-858-208-3810* where you can speak with an HR Consultant.

Action

If a Participant engages in any inappropriate behavior as defined herein, the Conference organizers may take action as deemed appropriate, including: a formal or informal warning to the offender, expulsion from the conference with no refund, barring from participation in future conferences or their organization, reporting the incident to the offender's local institution or funding agencies, or reporting the incident to local authorities or law enforcement. A response of "just joking" is not acceptable. If action is taken, an appeals process will be made available. There will be no retaliation against any Participant who brings a complaint or submits an incident report in good faith or who honestly assists in investigating such a complaint. All issues brought forth to the onsite HR Consultant during the course of a Conference will be immediately investigated.

Brief instructions for using Hopin

Upon joining Cosyne 2021 at https://hopin.com/events/cosyne-2021 you will arrive in the Reception area (as in the graphic above). Here you can review the Schedule, navigate through the different event areas using the left sidebar, and interact with fellow attendees using the right sidebar.

1. **Account Menu:** Your Profile serves as a virtual 'conference badge', and you can edit it from the drop-down menu here. We encourage you to complete your Profile by uploading a photo of yourself and filling in a brief bio; you may also add social media profile links.

2. **People tab:** Other participants in attendance are listed here. Clicking on an individual's profile gives you the option to contact them via direct message.

3. **Chat tab:** The Hopin platform provides multiple chat modes.

- The EVENT chat channel is available in all event areas, and you will find announcements posted here throughout the day.
- The STAGE chat channel – when in the Stage area use this chat channel to ask general questions of the organizers.
- The SLIDO chat channel – when in the Stage area use this chat channel to ask questions of presenters. The SLIDO chat allows other attendees to vote questions up and down (questions will be relayed by the Session Chair on screen).
- When in the SESSIONS and the POSTERS or EXPO area, there will be separate chats for each of these areas in addition to the EVENT chat.

4. This **"| >"** button minimizes the right sidebar.

5. **Stage area:** Talks will take place in the Stage area. When the talks are taking place, a red 'LIVE' icon will appear here.

6. **Sessions area:** You can create your own session (open or moderated) for discussions with colleagues and collaborators.

7. **Networking area:** When you click the 'Ready' button (and allow your audio and video to be shared), you will be randomly matched with another attendee in a one-on-one conversation for up to 10 minutes. Please note, you are able to leave the chat at any time, and do not need to wait the 10 minutes. Contact details will be exchanged if both people click the CONNECT button. Networking will be available throughout all the breaks during the event, and a red 'NOW' icon will appear here when available.

8. **Posters or Expo area:** You can preview posters in this area throughout the conference. Live poster presentations will be hosted on Gather.Town, please refer to the schedule for the link.

Hopin technical support and advice is available online https://support.hopin.to/en/collections/1945014-using-hopin-as-an-attendee.

In the event of any questions or issues on the day, please email *meeting@cosyne.org* or visit the Cosyne Virtual Help Desk in SESSIONS.

Brief instructions for using Gather.town

Joining the gathering

Cosyne poster sessions will take place on Gather.town. To access Cosyne 2021 space, use the following link: https://gather.town/app/nc4OmwdoklVsqGZx/cosyne. Only registered Cosyne attendees can attend poster sessions, so on gather.town you have to register using *the same email address that you used to register on Hopin.*

Requirements: a working microphone and camera; a web browser (Chrome or Firefox are recommended and known to work); headphones are strongly recommended to help prevent audio feedback. If you are using Windows or MacOS, you can also install a standalone gather.town app (https://gather.town/download).

Gather.town basics

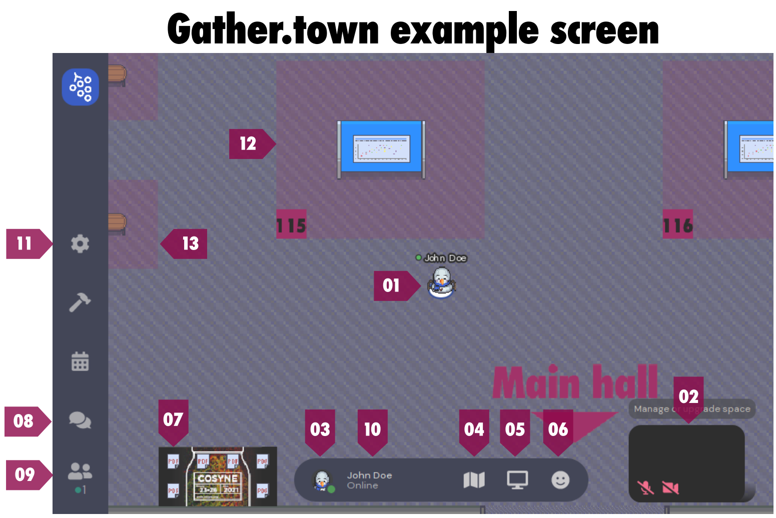

Gather.town is a video chat platform in which you can move your avatar (**01**) inside a 2D space. As you get closer to other avatars (users), your video and microphone (**02**) get connected and you will be able to chat. You can move around the space using the arrow keys. By walking your avatar you can have conversations with other nearby avatars (users). These conversations can include several people who are around you.

You can use several tools available at the bottom of your screen to change your avatar (**03**), see a minimap of the current room (**04**), share your screen (**05**), raise your hand (**06**), etc.

When you move close to an interactable object (poster, downloadable program **07**), there will be a notification at the bottom of your screen saying "Press x to to view the poster" or "Press x to download Cosyne program", for example.

There is a 'Chat' feature (**08**) that allows you to send messages to individual users (by clicking their name in the list of participants), nearby users, or everyone in the space.

You can locate other users in the space by clicking on their name in the list of participants (**09**) and using the 'Locate' feature.

If you wish to minimize the number of interactions, you can toggle 'Quiet Mode' on and off by pressing Ctrl-U (see also under **10**). In Quiet Mode you will only receive audio and video from people immediately next to you. If you turn your camera and microphone off, you will also minimize the number of 'live' interactions.

If you get lost or something goes wrong, please go to 'Settings' (**11**) and select 'respawn', which will reinitialize your position in Cosyne space.

Cosyne space

The main meeting space is composed of three rooms, Main hall and two Poster rooms (60 posters in each room per poster session). When you enter the space you will start in the Main hall. At the top of the Main hall there are four doors, two leading to 'Posters 001–060' and two leading to 'Posters 061–120'. The two poster rooms are also mutually connected by three doors. Each poster room contains sixty labeled rectangular poster booths (**12**). If you are a poster presenter, simply head to your designated poster booth in the beginning of the poster session.

In each Poster room there are also several private spaces where 2–4 people can have private conversations (**13**). These spaces contain tables in the middle, and only people inside the spaces, around the tables, can participate in the conversation.

The Cosyne 2021 program booklet will contain most up-to-date information about Cosyne space.

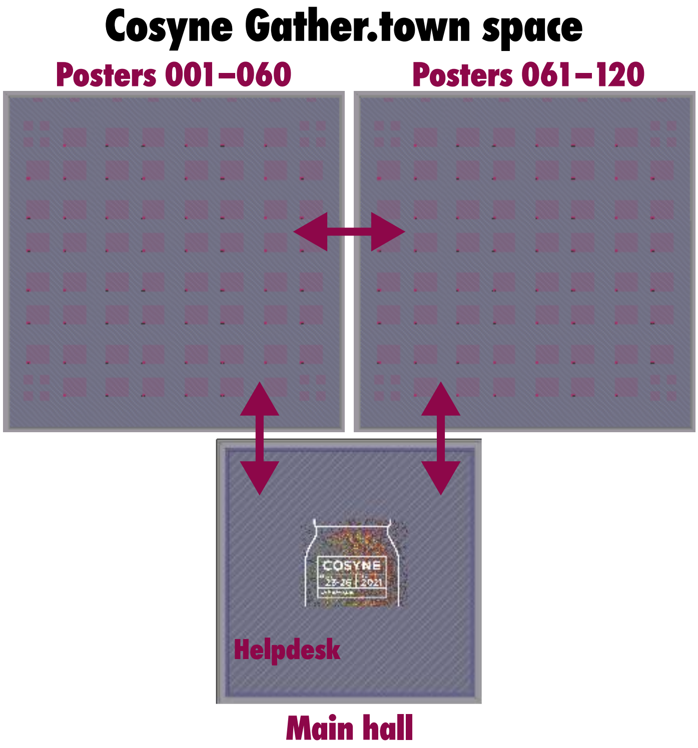

Poster booths

Each poster booth is designed as a 'private space' rectangle with poster number in its bottom left corner. When users approach the poster booth, they will see a poster thumbnail at the bottom of their screen, and when they press 'x' they will be able to browse the entire poster. Only after avatars enter the poster booth, they will engage in the conversation with all other avatars currently present in the poster booth. Since the poster booth is a private space, the conversation will include only people inside the poster booth.

Please be courteous when discussing posters. Only the poster presenter should share their screen when presenting the poster. When the presenter is not present in the poster booth, you can search for the presenter in the list of participants, and chat with her/him directly. Remember, however, that some presenters will not be available during all three parts of the poster session due to time difference.

Troubleshooting

If you encounter technical difficulties in gather.town, check if your browser permits using your camera and microphone. Refreshing the page and rejoining the gathering usually fixes most problems. Respawning to your initial position can help as well.

If your avatar gets ‘blocked’ by other avatars, you can turn on ‘Ghost mode’ by pressing Ctrl-G, move your avatar away, and then press G to turn the Ghost mode off.

There is also a ‘Helpdesk’ in the Main hall where you can find a moderator who might be able to troubleshoot your problems.

Program

Note: Institutions listed in the program are the primary affiliation of the first author. For the complete list, please consult the abstracts.

All times are EST (Eastern Standard Time, GMT-5)

Tuesday, 23 February

11:00a **Cosyne Tutorial: Recurrent Neural Networks for Neuroscience**

Session sponsored by the Simons Foundation

Lecturer: Kanaka Rajan (Mount Sinai)

11:00a Lecture on recurrent neural networks

12:45p Lecture of applications of recurrent neural networks to neuroscience

03:00p **Paths for leadership in promoting the careers of scientists**

This is a professional development and mentoring workshop consisting of an interactive discussion of three panelists with the COSYNE attendees on the paths for leadership in promoting the career of members of the underrepresented groups, including but not exclusive to women in neuroscience.

Organizers: Maria Geffen, Eugenia Chiappe

Panelists: Dr. Carmen Sandi (EPFL, director of the ALBA network), Dr. Emmeline Edwards (National Center for Complementary and Integrative Health, founder of World Women in Neuroscience), Dr. Susan Margulies (Georgia Institute of Technology and Emory University)

Wednesday, 24 February

08:00a **Poster Session 1a**

Session 1:

(Chair: Stephanie Palmer)

11:00a **Break**

Session 2:

(Chair: Tim Vogels)

01:00p **Simons Foundation Bridge to Independence Awards: Facilitating the transition from postdoc to junior faculty**

The Simons Foundation is invested in supporting the next generation of researchers. Our Bridge to Independence (BTI) Award programs promote talented early-career scientists by facilitating their transition to research independence and providing grant funding at the start of their professorships. We are offering such awards through two of our programs: the Simons Collaboration on the Global Brain (SCGB), and the Simons Foundation Autism Research Initiative (SFARI).

Speakers: Laura Long (SCGB), Alice Luo Clayton (SFARI)

Session 3:

(Chair: Christine Constantinople)

05:00p **Mini-symposium on the history of race and racism in science**

Race, it is commonly claimed, is a social construct of no biological significance. Yet, since the 1700s scientists have attempted to classify human populations into discrete groups. In this mini-symposium plus discussion, we will examine the history of race as a scientific concept and consider its abuses and misapplications. We will examine the cultural and historical context of scientific approaches to classifying humans into racial groups starting with Linneaus and Darwin. We will examine the abuses of race-based science manifest in the eugenics movement in the UK and USA and in Nazi Germany and explore how an academic discipline can become a justification for oppression and genocide. We will introduce the use of racial classifications in contemporary science and medicine through a historical lens. We will examine the utility of racial classification in identifying the genes that underlie human disease and understanding human history and variation in human behavior. Can we differentiate between race-based science and racist science? Is race-based science of potential use? Or, in order to avoid racism must we avoid race in science?

Speakers: David Gresham (New York University), Habon Issa (New York University)

Moderator: Wei Ji Ma (New York University)

07:00p **Poster Session 1c**

Thursday, 25 February

08:00a **Poster Session 2a**

Session 4:

(Chair: Athena Akrami)

11:00a **Break**

Session 5:

(Chair: Adrien Peyrache)

01:00p **Break**

Session 6:

(Chair: Ida Momennejad)

03:00p **Poster Session 2b**

05:00p **Expanding horizons: Diverse perspectives in computational neuroscience**

The aim of this event is to broaden participation from URMs at Cosyne and to highlight diverse areas of computational and systems neuroscience research. This event will feature short research talks from all the speakers.

Speakers: Tahra Eissa (University of Colorado at Boulder), Jay Gill (University of California, Los Angeles), Fernanda Matias (Universidade Federal de Alagoas), Ilenna Jones (University of Pennsylvania), Carlos Ponce (Washington University in St. Louis), Sridevi Sarma (Johns Hopkins University)

07:00p **Poster Session 2c**

Friday, 26 February

08:00a **Poster Session 3a**

Session 7:

(Chair: Anne-Marie Oswald)

10:00a Neural and computational mechanism of task switching
C. A. Duan, X. Li, C. Ma, C. Brody, Z. Zhang, J. Erlich, Chinese Academy of Sciences . . 35

10:20a Prospective representation of distal spatial goals in rodent Orbitofrontal cortex
R. Basu, R. Gebauer, T. Herfurth, S. Kolb, Z. Golipour, T. Tchumatchenko, H. Ito, Max Planck Institute for Brain Research 36

10:40a Slowly-evolving dopaminergic activity controls the timing of self-initiated movement
A. Hamilos, G. Spedicato, Y. Hong, F. Sun, Y. Li, J. Assad, Harvard Medical School 36

11:00a **Break**

Session 8:

(Chair: Demba Ba)

11:40a Expansion and contraction of resource allocation in sensory bottlenecks
L. R. Edmondson, A. Jimenez-Rodriguez, H. P. Saal, University of Sheffield 37

12:00n Dimensionality of neural activity in sensory cortex predicts activity propagation and behaviour
J. Rowland, R. Lees, M. Loidolt, T. van der Plas, J. Dehning, V. Priesemann, A. Packer, University of Oxford 37

12:20p Optimization for Real-World Environments Explains Biological Sound Localization
A. Francl, J. McDermott, Massachusetts Institute of Technology 38

12:40p Sinusoidal transformation of the visual field on the cortical surface
M. Sedigh-Sarvestani, D. Fitzpatrick, Max Planck Florida Institute of Neuroscience 38

01:00p **Break**

Session 9:

(Chair: Tim Hanks)

02:00p The geometry of the representation of decision variable and stimulus difficulty in the parietal cortex
G. Okazawa, C. Hatch, A. Mancoo, C. Machens, R. Kiani, New York University 39

02:20p Flexible identification of neural population dynamics underlying decision making
M. Genkin, C. Chandrasekaran, K. Shenoy, T. Engel, Cold Spring Harbor Laboratory . . . 39

02:40p Attention improves information flow between neuronal populations without changing the communication subspace
R. Srinath, D. A. Ruff, M. R. Cohen, University of Pittsburgh 40

03:00p **Poster Session 3b**

07:00p **Poster Session 3c**

Poster Session 1 Wednesday 24 February

Poster Session 2 Thursday 25 February

Poster Session 3 — Friday 26 February

Abstracts

Abstracts for talks appear first, in order of presentation; those for posters next, in order of poster session and board number. An index of all authors appears at the back.

T-1. Optimal information loading into working memory by non-normal attractor dynamics in PFC

Jake Stroud[1] JPS99@CAM.AC.UK
Kei Watanabe[2] KEIWATANABE1977@GMAIL.COM
Takafumi Suzuki[3] T.SUZUKI@NICT.GO.JP
Mark Stokes[4] MARK.STOKES@PSY.OX.AC.UK
Mate Lengyel[1,5] M.LENGYEL@ENG.CAM.AC.UK

[1]University of Cambridge
[2]Osaka University
[3]National Institute of Communication and Information Technology Osaka
[4]University of Oxford
[5]Department of Engineering

The paradigmatic neural circuit model of working memory is based on attractor dynamics giving rise to persistent activities. While these models have been used extensively to explain how information is maintained in working memory, how this information is first loaded into memory for subsequent maintenance has received far less attention. The prevailing dogma for information loading implicitly assumes that stimuli must drive neural activity directly into the desired pattern of persistent activity. Here, we show instead that the mechanism that loads information into working memory uses inputs that are in fact largely orthogonal to those that will eventually be expressed persistently. We first perform mathematical analyses of realistic non-normal attractor circuits to show that the most efficient mechanism for generating persistent activity is to use specifically tuned orthogonal inputs. We identify two neural signatures of this non-normal attractor mechanism: cross-temporal decoding and a novel measure of overlap of neural activity with optimal input vs. persistent dynamics. These measures distinguish between classical circuit models of working memory, based on either attractor or non-normal sequential dynamics, and our non-normal attractor model that unifies these two mechanisms and exhibits a combination of classical dynamics. Using function-optimized recurrent neural networks, we also show that these different dynamical regimes optimize distinct objectives: to produce target activities either at response time, throughout the delay period, or 'just-in-time' (before the response, at the end of the delay period). Finally, we analyze multiple datasets of prefrontal cortical recordings while monkeys perform memory-guided saccade tasks. Using our neural measures, we show that it exhibits clear signatures of non-normal attractor dynamics. These results also generalise to neural responses recorded in a task involving distractor stimuli. In summary, PFC working memory responses reveal non-normal attractor dynamics performing just-in-time computation in which optimal inputs are orthogonal to persistent modes.

T-2. Long-term spatial memory is maintained in the presence of continuous drift in hippocampal representations

Noa Sadeh[1] NONI.SDH@GMAIL.COM
Meytar Zemer[1] MEYTARZEMER@GMAIL.COM
Alon Rubin[1,2] ALON.RUBIN@GMAIL.COM
Yaniv Ziv[1] YANIV.ZIV@WEIZMANN.AC.IL

[1]Weizmann Institute of Science
[2]Neurobiology

A prevailing notion in memory research is that the persistence of a given memory depends on the stability of its coding by the neurons that were active during learning. Recent studies revealed that hippocampal place codes gradually change over time, even in unchanged familiar environments, calling into question the notion that stable neural codes underlie stable memories. However, several key questions are left open: (1) How do hippocampal place codes change under conditions that require stable long-term spatial memory? (2) To what extent does stable

performance in a spatial memory test rely on a stable hippocampal representation of the learnt environment? (3) Does re-training to navigate in a familiar environment weeks after learning reinstate the same representation that was present at the end of learning? We longitudinally imaged calcium dynamics and tracked large neuronal populations of the same neurons in the hippocampal CA1 of freely-behaving mice as they gradually learned to navigate to a reward zone in a radial arm maze, and then again 10 days or one month after learning. We found that CA1 representation of the maze was dynamic despite the need for- and successful expression of- long-term spatial memory. One month after learning, memory performance was correlated with place code quality but not with place code stability. Finally, additional training one month after learning led to rapid reinstatement of spatial memory performance and place-code quality, but the representation of the maze did not converge back to the representation that was established during learning. Together, our results suggest that spatial memory content is reliably preserved over weeks without relying on the long-term stability of hippocampal representations.

T-3. Structure and variability of optogenetic responses identify the operating regime of cortex

Agostina Palmigiano[1,2] AP3676@COLUMBIA.EDU
Francesco Fumarola[3,4] FRANCESCO.FUMAROLA@RIKEN.JP
Daniel Mossing[5] DAN.MOSSING@GMAIL.COM
Nataliya Kraynyukova[6] N.KRAYNYUKOVA@GMAIL.COM
Hillel Adesnik[5] HADESNIK@BERKELEY.EDU
Ken Miller[1] KDM2103@COLUMBIA.EDU

[1]Columbia University
[2]Center for Theoretical Neuroscience
[3]RIKEN Center for Brain Science
[4]Center For Brain Science
[5]University of California, Berkeley
[6]Max Planck Institute for Brain Research

What are regimes of operation of the cortical microcircuit? Modeling of recurrent networks with a single inhibitory type characterized two regimes: inhibition-stabilized, or not. In the former, steady-state inhibitory responses are "paradoxical": negative response to positive input. Here we expand this framework to study circuits with (i) multiple inhibitory subtype (PV, SOM, VIP) and (ii) heterogeneous connectivity and inputs; and (iii) fit these models to data. Studying low-dimensional models (one unit per cell type), we establish that a cell type's paradoxical response implies that the circuit without that cell type is unstable. Thus, observed paradoxical responses in PV cells imply PV cells stabilize the circuit. We fit the models to contrast responses of the four cell types, and find a common pattern of connectivity in models that fit the data as in experiments. Given this pattern, we show that responses of excitatory (E), PV, and SOM cells to SOM stimulation are negatively proportional to responses to VIP stimulation, with a common proportionality constant. We then consider high-dimensional models – many heterogeneous neurons of each cell type (inputs and connectivity determined by cell-type-specific or directed-cell-type-pair-specific normal distributions). We fit a known analytic form for the distribution of responses in these models to the observed cell-type-specific response distributions vs. contrast. Using mean field theory, we analytically characterize the model response distributions. We find a family of high-dimensional models that fit the data by generating model parameters from priors based on low-D fits, and accepting only solutions with a sufficiently low KL distance between data and mean field model distributions. We use random matrix theory to determine cell-type-specific response distributions to full or partial optogenetic stimulation. We show a "fractional paradoxical effect": the proportion of a cell type responding negatively changes non-monotonically with the fraction of cells of that type stimulated.

T-4. Impression learning: online predictive coding with synaptic plasticity

Colin Bredenberg[1,2] CJB617@NYU.EDU
Eero Simoncelli[1] EERO.SIMONCELLI@NYU.EDU
Cristina Savin[1,3] CS5360@NYU.EDU

[1]New York University
[2]Neural Science
[3]CNS, CDS

Early sensory areas in the brain are faced with a task analogous to the scientific process itself: given raw data, they must extract meaningful information about its underlying structure. This process is particularly difficult, because the true underlying structure of the data is never revealed, so representation learning must be largely

unsupervised. Framing this process in the language of Bayesian probabilities is tempting but difficult to connect to biology, because we still lack a satisfactory account of how the machinery of Bayesian inference and learning is implemented in neural circuits. Here, we provide a theoretical account of how learning to infer latent structure can be implemented in neural networks using local synaptic plasticity. To do this, we derive a learning algorithm in which synaptic plasticity is driven by a local error signal, computed by comparing stimulus-driven responses to internal model predictions (the network's 'impression' of the data). We associate these components with the basal and apical dendritic compartments of pyramidal neurons. Our solution builds on the Wake/Sleep algorithm (Dayan et al., 1995) by allowing learning to occur online, and capture temporal dependencies in continuous input streams. Compared to a traditional three-factor plasticity rule (Williams, 1992), it is substantially more stable and data-efficient, which allows it to be used for learning statistics of high-dimensional inputs. It is also flexible in that it is applicable to both rate-based and spiking-based neural activity, as well as different network architectures. More generally, our model provides a potential theoretical bridge from mechanistic accounts of synaptic plasticity to algorithmic descriptions of unsupervised probabilistic learning and inference.

T-5. Inhibitory feedback control of behavioral time scale synaptic plasticity induction in CA1

Sebi Rolotti[1,2] svr2111@columbia.edu
Mohsin Ahmed[1] msa2011@cumc.columbia.edu
Miklos Szoboszlay[1] ms5734@columbia.edu
Heike Blockus[1] hb2505@columbia.edu
Kevin Gonzalez[1] kg2685@columbia.edu
Fraser Sparks[1] fs2594@columbia.edu
Ana Sofia Solis Canales[1] assolisc@gmail.com
Attila Losonczy[1] al2856@columbia.edu

[1]Columbia University
[2]Department of Neuroscience

Pyramidal cells in area CA1 (CA1PCs) of the mammalian hippocampus develop spatially restricted firing fields as the organism explores an environment, forming a cognitive map. Clarifying the circuit mechanisms by which a given CA1PC forms a 'place field' (PF) is critical to a broader understanding of memory formation, and has been a topic of considerable interest.

Recent work has revealed that PF formation occurs through a novel form of plasticity, behavioral time scale synaptic plasticity (BTSP). Importantly, BTSP can be experimentally induced in individual CA1PCs via repeated current injection at a fixed location, allowing experimenters to probe PF dynamics and representational content. To date, the limitations of intracellular electrophysiology have constrained this avenue of inquiry to single cells over short timescales, leaving the relationship of BTSP to population coding and memory behavior undetermined. The induction and subsequent longitudinal tracking of PFs in arbitrary numbers of cells would therefore represent a major advance, clarifying the circuit mechanisms regulating BTSP and the relationship of PF formation in neural ensembles to learning and behavior.

Here we argue that recurrent inhibition in CA1 represents a circuit mechanism for limiting scaling of BTSP across the population. We present sparse-labeling and targeted-stimulation approaches that enable the first demonstrations of optogenetic PF induction in single cells. Pairing these optogenetic approaches with 2-photon calcium imaging, we characterize longer timescale dynamics of induced PFs to show they are functionally indistinguishable from spontaneously formed PFs. We find that these approaches scale to small numbers of stimulated cells but fail to induce PFs at higher densities. Importantly, we further show that local interneuron (IN) activity scales with stimulation density, and that manipulation of local INs to suppress recurrent inhibition raises the density limits on PF formation. Finally, we find that over-representing a location using these approaches can bias goal-directed behavior.

T-6. Dendritic Gated Network: A cerebellum-inspired alternative to backprop for learning

Sanjukta Krishnagopal[1,2] S.KRISHNAGOPAL@UCL.AC.UK
Eren Sezener*[3] ESEZENER@GOOGLE.COM
Agnieszka Grabska-Barwinska*[3] AGNIGB@GOOGLE.COM
Dimitar Kostadinov[1] D.KOSTADINOV@UCL.AC.UK
Maxime Beau[1] MAXIMEBEAUJEANROCH047@GMAIL.COM
David Budden[3] BUDDEN@GOOGLE.COM
Claudia Clopath[4] C.CLOPATH@IMPERIAL.AC.UK
Michael Hausser[1] M.HAUSSER@UCL.AC.UK
Peter Latham[2] PEL@GATSBY.UCL.AC.UK

[1]University College London
[2]Gatsby Computational Neuroscience Unit, UCL
[3]Google DeepMind
[4]Imperial College London

It is largely accepted that learning in the brain occurs as a consequence of changes in synaptic weights. However, the exact mechanism underlying these weight changes during learning is unknown. One commonly used mechanism - 'backpropagation' - has been incredibly successful for a variety of machine learning tasks. However, backprop is simply not biologically plausible. Here we introduce a Dendritic Gated Network (DGN), a variant of the Gated Linear Network architecture [Budden et. al. 2020], that offers a biologically plausible alternative to backprop through gating and local learning. The DGN architecture bears similarities to the cerebellum, where there is evidence for shaping of Purkinje cell responses by interneurons, which could be a natural candidate for 'gating'. In addition, the DGN is significantly more efficient than conventional artificial neural networks in terms of fast, online learning and retention of previously acquired tasks – important requirements for biological learning. We demonstrate our results on tasks involving resilience against catastrophic (long term) forgetting on MNIST, as well as vestibulo-ocular reflex (VOR), an eye movement task commonly studied in the neuroscience community, and show that the DGN matches state-of-the-art performance while displaying more efficient and natural learning. Properties of DGNs make it a preferred candidate for biological learning, bridging the fields of neuroscience and machine learning.

T-7. Predictable fluctuations in synaptic strength due to natural variation in presynaptic firing rate

Naixin Ren[1,2] NAIXIN.REN@UCONN.EDU
Abed Ghanbari[3] ABEDXGHANBARI@GMAIL.COM
Ian Stevenson[1] IAN.STEVENSON@UCONN.EDU

[1]University of Connecticut
[2]Department of Psychological Sciences
[3]Boston University

Many in vitro studies have shown how synaptic strength can vary with presynaptic spike timing due to short-term synaptic plasticity (STP). Synapses with short-term depression (STD), for instance, should be weaker when presynaptic firing rates are higher and stronger when rates are lower. Here we examine natural fluctuations in synaptic strength in vivo using large-scale spike recordings from the Allen Institute Neuropixels dataset [1]. Using cross-correlations between neuron pairs, we detect >9,000 putative excitatory synaptic connections between ~4 million neuron pairs. We then use an extension of the Generalized Linear Model to describe each cross-correlogram and to quantify fluctuations in synaptic efficacy over time. For individual connections, we find that the synaptic efficacy often varies substantially over time, and the efficacy fluctuations mirror fluctuations in the presynaptic firing rate (both measured on a minute timescale). The firing rates of the observed presynaptic neurons fluctuate over time by a factor of ~6 (median range), and the synaptic efficacies fluctuate between ~1% and 100% of their maximums (median range). We then fit a statistical model of STP to the observed pre- and post-synaptic spike trains – a generalized bilinear model (GBLM). Here we estimate a "modification function" that describes how the synaptic coupling changes following the occurrence of presynaptic spikes at a given inter-spike interval. We find that, for many connections, the GBLM accurately reproduces the observed efficacy fluctuations, suggesting that these fluctuations can be predicted based on the time-varying presynaptic firing rate and can be, at least partly, explained by a cumulative effect of STP. In vitro studies and theoretical models have described how STD could act as a gain control mechanism to balance the influence of multiple synaptic inputs [2]. Here we find evidence that gain control-like phenomena can occur due to natural, ongoing firing rate fluctuations in vivo.

T-8. Serotonin neurons modulate learning rate through uncertainty

Cooper Grossman[1,2] CGROSSM4@JHMI.EDU
Bilal Bari[1] BBARI1@JHMI.EDU
Jeremiah Cohen[1] JEREMIAH.COHEN@JHMI.EDU

[1]Johns Hopkins University School of Medicine
[2]Solomon H Snyder Department of Neuroscience

Regulating how fast to learn is critical for flexible behavior. Learning about the consequences of actions should be slow in stable environments, but accelerate when that environment changes. Recognizing stability and detecting change is difficult in environments with noisy relationships between actions and outcomes. Under these conditions, theories propose that uncertainty can be used to modulate learning rates ("meta-learning"). We show that mice behaving in a dynamic foraging task exhibit choice behavior that varied as a function of two forms of uncertainty estimated from a meta-learning model. The activity of dorsal raphe serotonin neurons tracked both types of uncertainty in the foraging task, as well as in a dynamic Pavlovian task. Reversible inhibition of serotonin neurons in the foraging task reproduced changes in learning predicted by a simulated lesion of meta-learning in the model. We thus provide a quantitative link between serotonin neuron activity, learning, and decision making.

T-9. Behavioral and dopaminergic signatures during chronic social defeat stress that predict resilience

Lindsay Willmore[1,2] WILLMORE@PRINCETON.EDU
Ilana Witten[1] IWITTEN@PRINCETON.EDU
Annegret Falkner[3] AFALKNER@PRINCETON.EDU

[1]Princeton University
[2]Neuroscience Institute
[3]Princeton Neuroscience Institute, Princeton University

Chronic stress can have lasting adverse consequences in some individuals, yet others exhibit resilience. In the chronic social defeat stress (CSDS) model, socially defeated mice are characterized as stress-susceptible or resilient based on the degree of social avoidance after defeat. While many previous studies have focused on identifying neural and molecular correlates of resilience following defeat, few have looked at how resilience may be predicted by behavior or neural correlates expressed during the stress experience itself. To address this gap, here we combine behavior tracking with bulk calcium recordings in the dopaminergic system across the ten days of social defeat, as mice undergo repeated attacks from aggressor mice. Using supervised behavioral quantification, we find that while susceptible and resilient mice both receive similar levels of attack, susceptible mice employ a more consistent behavioral strategy in response (retreating) and resilient mice show diverse behavioral strategies, which include more active resistance (fighting back). These behavioral differences are reflected by differences in the response of dopaminergic neurons, where, across animals, the level of active behavioral resistance to attack is correlated with the amount of dopamine signal at the attack onset. Taken together, our behavioral and neural findings suggest that susceptible mice, through consistent retreats and consistent dopamine dips at aggressor encounters, develop lasting negative associations with the aggressor strain, while resilience may be a product of either increased active resistance behavior or more variable dopamine responses to attack, both suggestive of diminished avoidance learning. We are currently using closed-loop behavior-triggered optogenetic stimulation to test whether manipulating dopaminergic activity at fighting onset is sufficient to increase fighting back behavior and ultimately promote resilience following defeat.

T-10. Choice-selective sequences dominate in cortical relative to thalamic inputs to nucleus accumbens, providing a potential substrate for credit assignment

Avinash Baidya[1,2] AAVINASH@UCDAVIS.EDU
Nathan Parker[3] NPARKER@ROCKEFELLER.EDU
Julia Cox[4] JULIAMC@PRINCETON.EDU
Ilana Witten[5] IWITTEN@PRINCETON.EDU
Mark Goldman[1] MSGOLDMAN@UCDAVIS.EDU

[1]University of California, Davis
[2]Department of Physics and Astronomy
[3]The Rockefeller University
[4]Princeton Neuroscience Institute, Princeton University

[5]Princeton University

A central question in reinforcement learning is how actions and outcomes separated in time become associated with each other. Multiple lines of evidence have implicated the nucleus accumbens, a part of the ventral striatum that receives cortical and thalamic inputs as well as dopaminergic signals associated with reward prediction errors, in learning action-outcome associations. Here we image the activity of thalamic and cortical neurons projecting directly to the nucleus accumbens during a two-armed bandit task in which an animal must determine the lever corresponding to the higher probability of reward. Prelimbic cortical inputs preferentially encoded animals' recent actions (choices) on a given trial, whereas thalamic neurons preferentially encoded external cues (stimuli). The cortical input population exhibited choice-selective sequences of neural firing that bridged a delay period between action (choice) and outcome (reward). We construct a set of computational models of the underlying cortico-striatal circuitry to demonstrate how such choice-selective sequences can create the precisely timed reward prediction error seen in dopamine neurons and support the neural implementation of temporal difference learning, a powerful reinforcement learning algorithm for associating actions and outcomes separated in time. By optogenetically manipulating both the cortical and thalamic input neurons we tested and confirmed core predictions of our circuit models. Thus, by combining experiments and computational modeling, we provide a neural circuit solution to the credit assignment problem and suggest a computational role for choice-selective sequences of neural activity.

T-11. Gated feedforward inhibition in the frontal cortex releases goal-directed action

Jae-Hyun Kim[1,2] ZKJH@KAIST.AC.KR
Dong-Hyun Ma[1] MDH8685@KAIST.AC.KR
Eunji Jung[1] EUNJI_JUNG@KAIST.AC.KR
Ilsong Choi[1] MARSHMALLOW@KAIST.AC.KR
Seung-Hee Lee[1] SHLEE1@KAIST.AC.KR

[1]Korea Advanced Institute of Science and Technology
[2]Biological Sciences

Perception of sensory information triggers precise motor actions in animals performing perceptual tasks. However, it is still unclear how cortical circuits permit appropriate responses to relevant sensory stimuli while inhibiting impulsive actions. Here, we identified a visual-to-motor inhibitory circuit in the cingulate cortex (Cg) that releases goal-directed action in mice performing the visual Go/No-go task. Using in vivo multichannel recordings, we identified two populations of neurons in the Cg of task-performing mice: one showing visual responses with gain modulation and the other showing ramp-to-threshold activity before licking. Interestingly, amplitude in the visual responses and timing in the ramp-down-to-threshold activity before licking showed a significant correlation with lick latency. Furthermore, the optogenetic activation of Cg sensory neurons that receive direct inputs from the visual cortex (VC) was sufficient to suppress neighboring Cg neurons and generate licking action even in the absence of visual stimuli. Notably, neighboring neurons, which were neither sensory nor motor, showed a gating signal that inhibited sensory-to-motor suppression, during which their activity decreased in correlation with licking. Our findings illustrate how neural circuits in the Cg transform salient visual cues into precise motor actions. The increase of sensory signals suppressed sustained motor activity in the Cg. Neighboring neurons gate this sensory-to-motor suppression and, in turn, released licking action. Together, our data demonstrate that visual inputs to the frontal cortex exert gated feedforward inhibition of motor neurons to trigger goal-directed actions. Importantly, Cg resolves conflicts between the inhibition and the execution of goal-directed action by gating the sensory-induced suppression of the sustained motor activity that restrains impatient responses.

T-12. Distinctive signatures of information seeking during changes of mind in controllable environments

Marion Rouault[1,2] MARION.ROUAULT@GMAIL.COM
Aurelien Weiss[1] AURELIENWEISS@GMAIL.COM
Jun Seok Lee[1] JUN.SEOK.LEE@ENS.FR
Jules Boute[3] JULESBOUTE@HOTMAIL.FR
Jan Drugowitsch[4,5] JAN_DRUGOWITSCH@HMS.HARVARD.EDU
Valerian Chambon[1] VALERIAN.CHAMBON@ENS.FR
Valentin Wyart[1] VALENTIN.WYART@ENS.FR

[1]Ecole Normale Superieure
[2]Department of Cognitive Studies

[3]Paris-Saclay Institute of Neuroscience
[4]Harvard Medical School
[5]Neurobiology

The ability to flexibly adapt our beliefs and behaviour in light of new evidence is a hallmark of human cognition. Using perceptual decision-making as a model system, previous re- search has focused on situations in which participants are passive observers adapting their behaviour to sensory evidence. However, outside the laboratory, humans typically exert active control over their environment by seeking information to guide their behaviour. Here, we examined how being in control over decision evidence influences changes-of-mind and their associated confidence. Across five variants of a novel cognitive task for ensuring generalisability of our findings (N=113), we found that changes-of-mind required more evidence against current beliefs in controllable environments. These changes of mind were also associated with reduced confidence, but were nevertheless more likely to be confirmed on the next trial. Computational modelling explained these behavioural effects through more stable beliefs, rather than a reduced sensitivity to incoming evidence in controllable environments. Multimodal neurophysiological recordings showed that changes-of-mind occurring in controllable environments were preceded by a stronger involvement of a dorsal attentional network in magnetoencephalography, and followed by an increased pupil-linked arousal during decision evaluation. By bridging across historically distinct fields of decision research, these findings characterise specific signatures of human information seeking in controllable environments.

T-13. Serial time-multiplexed incorporation of evidence to make two decisions about one object

Danique Jeurissen[1] JJJ2132@COLUMBIA.EDU
Anne Loffler[1] AL3928@COLUMBIA.EDU
Yul Kang[2,3] YUL.KANG.ON@GMAIL.COM
Ariel Zylberberg[4] ARIEL.ZYLBERBERG@GMAIL.COM
Daniel M. Wolpert[1] WOLPERT@COLUMBIA.EDU
Michael N. Shadlen[5] MS4497@COLUMBIA.EDU

[1]Zuckerman Mind Brain Behavior Institute, Columbia University
[2]University of Cambridge
[3]Department of Engineering
[4]Department of Brain and Cognitive Sciences, University of Rochester, Rochester; Columbia University
[5]Zuckerman Mind Brain Behavior Institute, Columbia Uni.; Kavli Institute; Howard Hughes Medical Inst.

The brain is capable of processing several streams of information that bear on different aspects of the same problem. Here we address the problem of making two decisions about one object, by studying difficult perceptual decisions about the color and motion of a dynamic random dot display. We show that the accuracy of one decision is unaffected by the difficulty of the other decision when the double-decision is reported with a single hand movement to one of four targets. However, the response times (RT) reveal that the two decisions do not form simultaneously. Instead, the RT is determined by the sum of the decision times for motion and color, indicating a serial process. We then test whether the seriality is due to the use of a single effector. In a bimanual task participants report their two decisions with two hands, which might facilitate parallel integration of the two streams of evidence. However, we again find that the final RT reflects the sum of the motion and color decision times. When examining the first of the two movements in the bimanual task we find that the RT of the first completed decision is modulated by the difficulty of the second decision. This suggests that participants do not solve one decision before the other, but instead, use time-multiplexing for the two decisions. We propose a computational model where two streams of evidence are integrated in a serial, time-multiplexed manner due to a central bottleneck that precludes updating more than one decision at a time. We suggest that this bottleneck is responsible for the long timescales of many cognitive operations framed as decisions.

T-14. Wiring logic of the rodent olfactory system revealed by sequencing of single cell projections

Yushu Chen[1,2] YUCHEN@CSHL.EDU
Xiaoyin Chen[1] XICHEN@CSHL.EDU
Batuhan Baserdem[3] BASERDEM.BATUHAN@GMAIL.COM
Justus Kebschull[4] JUSTUS@KEBSCHULL.ME
Huiqing Zhan[1] HZHAN@CSHL.EDU
Yan Li[1] YALI@CSHL.EDU
Martin Davis[1] DAVISM@CSHL.EDU
Anthony Zador[1] ZADOR@CSHL.EDU
Alexei Koulakov[1] KOULAKOVALEXEI@GMAIL.COM
Dinu Albeanu[1] ALBEANU@CSHL.EDU

[1]Cold Spring Harbor Laboratory
[2]Neuroscience
[3]Stony Brook University
[4]Stanford University

The olfactory system relies on hundreds of odorant receptors to recognize numerous volatile compounds. However, the neuronal substrates of odor perception remain elusive, partially due to lack of insight into how the brain sorts the inputs from the sensory periphery. In contrast to the structured long-range connectivity and spatial representations in other sensory modalities, previous axonal tracing studies found that olfactory bulb (OB) outputs, the mitral and tufted cells, project in a highly distributed and seemingly random fashion to their largest cortical target, the piriform cortex (PC). However, previous studies were limited by the small number of neurons whose projections were traced, which may obscure any potential spatial structure in OB output. To investigate the wiring logic of the olfactory system, here we analyze the projections of 5,309 OB output neurons and 30,433 PC output neurons at single-cell resolution. We employ highly multiplexed barcode sequencing-based mapping techniques (MAPseq and BARseq), which enable discriminating different types of OB output neurons by their soma locations and brain-wide projections. We identify structured OB (mitral cell) projection modules, reproducible across mice, which tile the anterior-posterior (A-P) axis of the piriform cortex and co-innervate distinct brain regions. In addition, we find characteristic sets of PC output neurons which project to different brain targets, and are enriched at specific A-P locations in the piriform cortex. Drawing from these results, an organizing principle emerges: matched input-output, parallel circuit motifs span the A-P axis of the piriform cortex and co-innervate specific sets of OB target areas. These results challenge the canonical framework of piriform cortex as a random network, at the core of current understanding of olfactory processing. They encourage further investigation into the function of these neural circuits, and the formulation of novel computational models to account for the logic of information flow in the olfactory system.

T-15. Bilateral sensory signals for odor source localization in freely-moving mice

Kevin Bolding[1,2] K.K.BOLDING@GMAIL.COM
Ian Davison[1,3] IDAVISON@BU.EDU
Dan Leman[4] DPLEMAN@BU.EDU
Jiayue Tai[1] JTAI@BU.EDU

[1]Boston University
[2]Biology
[3]Department of Biology
[4]Brandeis University

During sensory-guided navigation, animals constantly refine their ongoing movement through a series of dynamic, iterative sensory-motor algorithms. In natural contexts, odors form key navigational cues that signal food sources and social partners, offering an ethologically relevant window on motivated sensory search. While odor-evoked activity has been studied intensively in head-fixed an- imals and decision-making tasks, little is known about the nature of the dynamic sensory signals that guide freely moving animals during active sampling of their environment. Such data is critical to resolve whether animals navigate using comparison of signals across successive discrete ‘sniff’ samples, using instantaneous ‘stereo’ comparison across hemispheres, or employ both under different conditions. To overcome the challenges of measuring bilateral odor responses in unrestrained animals, we devel- oped new miniaturized microscopy tools for large-scale visualization of neural activity, and used them to image both hemispheres of the main olfactory bulb in mice actively exploring odor sources in an open arena. We find that in the absence of airflow, olfactory information was detectable only within a restricted area of ~10 cm surrounding the odor source. Increasing proximity to the source progressively activated additional input

channels ('glomeruli'), revealing that spatial information is encoded in the progressive, concentration-dependent recruitment of receptor pathways of varying affinity. Directional tuning emerges in a subset of glomeruli at close proximity, where glomeruli in both left and right hemi- spheres displayed a bias in directional preference for stimuli nearest the corresponding naris. These data suggest that animals may employ multiple strategies to localize odor sources during free explora- tion, initially comparing the degree of glomerular recruitment across time during early approach phases, and ultimately reading out a bilateral direction code at close proximity.

T-16. Rat ACC continuously signals decision variables in a patch foraging task

Gary Kane[1,2] GAKANE@BU.EDU
Nathaniel Daw[3] NDAW@PRINCETON.EDU
Morgan James[4] MHJ32@CA.RUTGERS.EDU
Amitai Shenhav[5] AMITAI_SHENHAV@BROWN.EDU
Gary Aston-Jones[4] GSA35@CA.RUTGERS.EDU
Jonathan Cohen[3] JDC@PRINCETON.EDU

[1]Boston University
[2]Department of Psychological and Brain Sciences
[3]Princeton University
[4]Rutgers University
[5]Brown University

In patch foraging tasks, animals must decide to stick with a depleting resource or to leave it in search of a potentially better source. In such tasks, animals consistently follow the general predictions of optimal foraging theory (the Marginal Value Theorem; MVT): to leave a patch when the reward rate in the current patch depletes to the average reward rate across patches. Prior studies implicate an important role for the anterior cingulate cortex (ACC) in foraging decisions based on MVT: within single trials, ACC activity increases immediately preceding foraging decisions, and across trials, these dynamics are modulated as the value of staying in the patch depletes to the average reward rate. To examine what drives these activity patterns, we developed a leaky accumulator model based on the MVT that generates estimates of decision variables not only across trials, but also their dynamics within trials. We tested these predictions against ACC activity recorded from rats performing a patch foraging task, and found that the model-predicted changes in MVT-derived decision variables within and across trials closely matched rat ACC activity, suggesting that the ACC continuously signals MVT-derived decision variables. Next, we pharmacologically inactivated ACC to test the contribution of these signals to foraging decisions. Rats stayed in patches longer, but they still followed qualitative predictions of the MVT, suggesting that foraging decision variables represented in the ACC are not used to guide decision-making, but rather for a more general function such as monitoring ongoing task performance.

T-17. Transformation of velocity signals from egocentric to allocentric coordinates

Jenny Lu[1,2] JENNY_LU@HMS.HARVARD.EDU
Elena Westeinde[1] EWESTEINDE@G.HARVARD.EDU
Lydia Hamburg[3] LYDIAHAM@STANFORD.EDU
Shaul Druckmann[3] SHAULD@STANFORD.EDU
Rachel Wilson[1] RACHEL_WILSON@HMS.HARVARD.EDU

[1]Harvard Medical School
[2]Neurobiology
[3]Stanford University

Accurate navigation requires the brain to keep track of body velocity in allocentric coordinates. However, the motor commands and proprioceptive feedback signals that are used to estimate body velocity are represented in egocentric coordinates. Velocity information must therefore be transformed from egocentric to allocentric space. How the brain solves this problem is an open question. It is known that the insect brain contains "compass neurons," which form a topographic map of heading in allocentric coordinates (Seelig and Jayaraman, 2015). Locomotion involves both forward movement (aligned with heading) and sideways movement (orthogonal to heading). We show how heading information is combined with locomotion information. We identify two cell types in the Drosophila brain, called PFNd and PFNv, which encode both forward and sideways velocity. Each population responds in a graded manner to speed increases in its preferred direction. Each population also contains a topographic map of heading inherited from compass neurons. Output synapses from PFNd and PFNv neurons converge onto the

same cell type (hDeltaB). The synaptic outputs of forward- and backward-preferring neurons are precisely offset to reflect their different direction tuning. Moreover, the output synapses from neurons with opposite sideways velocity preferences are offset in opposite directions in the heading map, so that sideways movement becomes equivalent to forward movement with a heading shift. As a result, egocentric velocity representations are transformed into a topographic map of allocentric velocity. We formalize these notions in a computational model that draws its connectivity directly from the Drosophila central complex connectome, with egocentric velocity tuning profiles taken from our physiology data. By integrating these allocentric velocity signals over time, the brain should be able to track the animal's path through the environment.

T-18. Deconstructing meaning from human language at single-neuronal scales via deep learning models

Jing Cai[1] JCAI5@MGH.HARVARD.EDU
Yonatan Belinkov[2] BELINKOV@TECHNION.AC.IL
Mohsen Jamali[1,3] MJAMALI@MGH.HARVARD.EDU
Benjamin Grannan[1,3] BGRANNAN@PARTNERS.ORG
Ziv Williams[1,3] ZWILLIAMS@MGH.HARVARD.EDU

[1]Massachusetts General Hospital
[2]Technion - Israel Institute of Technology
[3]Harvard Medical School

Human language is among the most complicated cognitive processes that distinguishes humans from all other species. Through the basic combination of words and sentence constructions, humans are able to convey extraordinarily rich and nuanced meanings. Yet, despite its importance, the precise cellular and high temporal-resolution computations by which humans derive meaning from language have remained a major challenge to understand. Here, we used a rare opportunity to perform single-neuronal recordings from the human language-dominant prefrontal cortex of participants undergoing planned intraoperative neurophysiology, as participants listened to sentences across varying topics. By stably tracking the activity patterns of hundreds of neurons and by using advanced natural language processing deep learning models, we demonstrate that these cells carried a remarkably wide-range of linguistic information that reflected lexico-semantic (word meaning), combinatorial (syntactic) and working memory processes. We show how these single-neuronal 'building blocks' collectively reflected high-level linguistic representations and were reliably predictive of upcoming word meanings prior to utterance. We also show how the activities of these cell ensembles preferentially mapped onto 'top' network layers described by the deep learning models, suggesting their role in the integration of sentence-level information and deriving inferred meaning. Together, our findings reveal remarkably detailed linguistic processes and meaning prediction representations in the human prefrontal cortex at a single-cellular level and demonstrate the power of combining deep learning models with high spatiotemporal resolution neuronal recordings to better understand human cognitive processes such as language.

T-19. Ensembles of basal ganglia neurons drive, encode, and regulate vocal variability in the zebra finch

Jonnathan Singh Alvarado[1,2] JS570@DUKE.EDU
Jack Goffinet[1] JACKGOFFINET@GMAIL.COM
Richard Mooney[1] MOONEY@NEURO.DUKE.EDU

[1]Duke University
[2]Neurobiology

Skilled movements are typically variable during practice yet highly stereotyped during performance. Such behavioral variability is necessary for exploration in a reinforcement learning context, but it can be detrimental when precision is required. How ensembles of neurons encode and dynamically regulate motor variability across practice and performance states remains unknown. Male zebra finches sing more variable songs when practicing alone and highly stereotyped songs when performing to a female, providing a powerful system to explore how neural ensembles encode and regulate motor variability. Here we show that ensembles of spiny neurons (SNs) in the basal ganglia drive and encode vocal variability during practice, and that SN activity is strongly suppressed to enable stereotyped song performance in the presence of a female. During vocal practice, microendoscopic imaging reveals that calcium signals in ensembles of SNs are specific to song, premotor in nature, insensitive to auditory errors and highly variable relative to their cortical premotor afferents. In contrast, SN activity is strongly suppressed when the male sings to a female, and optogenetically inhibiting SNs strongly reduces vocal variability during solo practice, indicating that SN activity is causally linked to vocal variability. A variational autoencoder

(VAE) model trained to jointly encode neural activity and song spectrograms into a shared low-dimensional latent space reveals a coding scheme whereby specific patterns of SN activity map onto distinct spectral variants of syllables during solo practice. Finally, we establish that adrenergic signaling in the basal ganglia regulates vocal variability by directly suppressing SN activity. Thus, SN ensembles encode and drive vocal exploration during practice, and the social context-dependent regulation of SN activity enables stereotyped and highly precise vocal performance.

T-20. Driving neural co-firing with low-frequency stimulation restores dexterity after stroke in NHPs

Preeya Khanna[1,2] PREEYA.KHANNA@UCSF.EDU
Douglas Totten[1] DTOTTEN@UCDAVIS.EDU
Lisa Novik[3] LMNOVIK@UCDAVIS.EDU
Jeffrey Roberts[3] JAROBERTS@UCDAVIS.EDU
Karunesh Ganguly[1] KARUNESH.GANGULY@UCSF.EDU

[1]University of California, San Francisco
[2]Neurology
[3]California National Primate Research Center, Davis

Electrical stimulation is a promising tool for modulating brain networks. However, it is unclear how stimulation interacts with neural patterns underlying behavior. Neural patterns that produce behavior are thought to be driven by an interplay between “internal dynamics”, the property that neural activity transitions reliably in time without apparent influence from external sources1, and external inputs. Such internal dynamics are likely the result of network connectivity that produces sustained, time-varying neural patterns without need for external drive. How then can electrical stimulation, which is not sensitive to the state of ongoing internal dynamics, be used in a therapeutic manner to reliably assist neural processing and improve function?

Here, we tested how low-frequency epidural alternating current stimulation (ACS) in non-human primates (NHPs) recovering from sensorimotor stroke interacted with task-related activity in perilesional cortex (PLC) and affected grasping. We found a significant improvement in dexterity with ACS. We also found that PLC neurons entrained to the ACS waveform and across the recorded population, exhibited a common preferred ACS phase associated with increased firing. Further, ACS increased co-firing within task-related neural activity patterns in a phase dependent manner. Trials where movements began at the preferred phase of the ACS waveform exhibited stronger increases in co-firing and greater improvement in dexterity. These results suggest increased co-firing in task-related activity may enhance behavior in damaged networks.

A possible mechanism by which increased co-firing may improve behavior is by enhancing propagation of signals through damaged networks2. In support of this interpretation, we developed a neural network model of the recovering PLC and found that when ACS was applied to the model, phase-dependent increases in co-firing were observed and predicted better propagation of signals through impaired parts of the network. Overall, increases in population co-firing may enhance signal propagation in damaged networks and improve behavior.

T-21. Neural and computational mechanism of task switching

C. Ann Duan[1,2] ANNDUAN2@GMAIL.COM
Xiangci Li[3] LIXIANGCI8@GMAIL.COM
Ce Ma[2] CM4913@NYU.EDU
Carlos Brody[4] BRODY@PRINCETON.EDU
Zheng Zhang[5] ZZ@NYU.EDU
Jeffrey Erlich[2] JERLICH@NYU.EDU

[1]Chinese Academy of Sciences
[2]New York University Shanghai
[3]University of Texas at Dallas
[4]Princeton Neuroscience Institute
[5]AWS Shanghai AI Lab

Behavioral flexibility is a core component of higher cognition. One prominent feature of this flexibility is that there is a “switch cost”, an increase in errors and reaction time, when switching from one task to another based on the relevant goal at hand. It has been hypothesized that a major source of task switch cost could be 1) the residual interference from the previous task (“task-set inertia”); or 2) the difficulty in setting up the new task (“task-set reconfiguration”)1. Despite behavioral results that are broadly consistent with both of these theories, direct neurophysiological evidence in support of either theory is lacking. To study the mechanism for task switching, we

trained rats and recurrent neural networks (RNN) to perform a task-switching paradigm using similar procedures. Both rats and RNNs developed switch cost, which diminished with more training. These behavioral data suggest that switch cost likely results from a lack of experience on switching, and the underlying mechanisms can be studied in artificial networks as a complementary approach. Analysis of activity from real and artificial neurons of well-trained rats and RNNs provided correlational evidence that the cost of switching is largely due to the residual memory of the previous task, which interferes with the current task. Silencing the "residual memory" in RNNs during the inter-trial interval (ITI) while preserving the input activity to RNNs on the current trial eliminated switch cost. Additionally, inserting a memory of the opposite task before a non-switch trial created a cost even in the absence of task switching. These experiments provide in silico causal evidence for the residual memory interference hypothesis. Together, our results provide insights to the source of task switch cost and demonstrate how ongoing neural activity, shaped by recent experience, can influence new tasks at hand.

T-22. Prospective representation of distal spatial goals in rodent Orbitofrontal cortex

Raunak Basu[1] RAUNAKBASU.IITKGP@GMAIL.COM
Robert Gebauer[1] ROBERT.GEBAUER@BRAIN.MPG.DE
Tim Herfurth[1] TIMHERFURTH87@GMAIL.COM
Simon Kolb[1] SIMONKOLB@GOOGLEMAIL.COM
Zahra Golipour[1] ZAHRA.GOLIPOUR@BRAIN.MPG.DE
Tatjana Tchumatchenko[2] TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE
Hiroshi Ito[1] HIROSHI.ITO@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]MPI for Brain Research

Planning of goal-directed actions requires relationships between current and future behavioral states. In spatial navigation, it has been thought that the brain needs a representation of the next goal at the beginning of the journey. However, while neurons in the hippocampus and parahippocampal regions provide estimates of the animal's positions as well as nearby trajectories, a specific neural code of the remote navigational goals remains to be identified. Here we report that neurons in the rat orbitofrontal cortex (OFC) forms a spatial map of the environment, discriminating positions on the maze during navigation. The activity of OFC neurons encodes animal's positions during reward consumption, but before the onset of journey, these neurons exhibit a representational transition from the animal's current position to the subsequent goal, which is kept during the entire journey until the animal reaches the destination. Further analysis revealed that this persistent goal representation is maintained by destination specific neural trajectories that are predefined at navigation onset and embedded within underlying neural dynamics which evolve along animal's trajectory towards the goal. Finally, optogenetic perturbation of OFC neurons, specifically at the navigation onset, caused animals to navigate to an incorrect destination, pointing to OFC as part of the brain's internal map representing the animal's decision of navigational goals.

T-23. Slowly-evolving dopaminergic activity controls the timing of self-initiated movement

Allison Hamilos[1,2] AHAMILOS@G.HARVARD.EDU
Giulia Spedicato[1] GIULIASPEDICATO@GMAIL.COM
Ye Hong[1] HONGYE161@MAILS.UCAS.AC.CN
Fangmiao Sun[3] SUNFANGMIAO@GMAIL.COM
Yulong Li[3] YULONG@GMAIL.COM
John Assad[1] JASSAD@HMS.HARVARD.EDU

[1]Harvard Medical School
[2]Neurobiology
[3]Peking University School of Life Sciences

Parkinson's disease has long suggested dopamine is vital for movement initiation, yet there is currently scant evidence linking the endogenous dopaminergic system to this process. For example, although dopamine neurons show brief bursts of activity before movements, optogenetic stimulations mimicking such bursts do not elicit movements (Coddington & Dudman, Neuron, 2019), and just how this system participates in movement initiation remains elusive.

We took a new approach to this question. Previous studies examined movements made in rapid response to external stimuli—but most natural movements don't require explicit triggering. One such category is self-timed movements, which occur without prompting after an elapsed interval. In fact, decades-old pharmacological studies

suggest dopamine somehow influences the timing of such movements. We reasoned self-timing could be key to illuminating the question of the endogenous dopamine system's role in movement onset–which has been missed under "standard," rapid-reaction paradigms.

We recorded midbrain dopaminergic activity and striatal dopamine release as mice executed a self-timed movement task. We discovered a slow dopaminergic "ramp-up" present while the animals were measuring time before movement. Remarkably, the ramp's slope reflected the varying timing from trial-to-trial: steep ramping preceded relatively early movements and vice versa, allowing us to precisely predict the upcoming movement time on single trials—even many seconds before the actual movement. Employing several complementary analytical techniques, we found a strong correlation between dopaminergic activity and movement timing, but the relationship was also causal: critically, optogenetic activation of dopamine neurons shifted the distribution of movement-times earlier, whereas inhibition caused late-shifting. Slowly varying dopaminergic signals have been described in certain complex behavioral settings, but our results mark the first time these have been concretely linked to a precise motor output. These findings provide the first clear view of how dopamine neurons influence movement initiation, an important, longstanding question in neuroscience.

T-24. Expansion and contraction of resource allocation in sensory bottlenecks

Laura R Edmondson[1,2] LREDMONDSON1@SHEFFIELD.AC.UK
Alejandro Jimenez-Rodriguez[1] A.JIMENEZ-RODRIGUEZ@SHEFFIELD.AC.UK
Hannes P. Saal[1] H.SAAL@SHEFFIELD.AC.UK

[1]University of Sheffield
[2]Department of Psychology

In many sensory systems, information from receptors is projected to subsequent processing stages in a topographic fashion, such as retinotopy in vision. However, these projections are typically not scaled proportionally to the size of the input regions and the representations of some regions will expand, while others contract. For example, in vision, the area representing the fovea is expanded in primary visual cortex, and in touch, the highly innervated fingertips occupy an outsized representation in primary somatosensory cortex. What principles drive the allocation of cortical resources, and how should receptor density, e.g. the high innervation of the fovea or the fingertips, and stimulus statistics, e.g. the much higher contact frequencies on the fingertips, contribute? Here, building on prior work in efficient coding, we solve this problem both analytically and numerically in linear second-order models that maximize information transmission through decorrelation. We introduce a sensory bottleneck to impose strict constraints on resource allocation and derive the optimal neural allocation for different bottleneck widths while varying the receptor density and stimulus statistics on the sensory sheet. Surprisingly, we find that bottleneck width is a crucial factor in the resulting resource allocation, and changing this one parameter can induce both expansion and contraction of individual input regions. Both the density of receptors and the stimulus statistics affect the resulting allocation and determine the convergence point for wider bottlenecks. Furthermore, we show a close match between the predicted optimal allocations and empirical somatosensory cortical representations in a well-studied model system, the star-nosed mole. Overall, our results suggest that the presence and strength of cortical magnification that is empirically observed in many sensory systems might depend on the amount of resources made available.

T-25. Dimensionality of neural activity in sensory cortex predicts activity propagation and behaviour

James Rowland[1,2] JROWLANDM@GMAIL.COM
Robert Lees[1] ROBERT.LEES@DPAG.OX.AC.UK
Matthias Loidolt[3] MATTHIAS.LOIDOLT@DS.MPG.DE
Thijs van der Plas[1] VDPLASTHIJS@GMAIL.COM
Jonas Dehning[3] JONAS.DEHNING@DS.MPG.DE
Viola Priesemann[3] VIOLA.PRIESEMANN@DS.MPG.DE
Adam Packer[1] ADAMPACKER@GMAIL.COM

[1]University of Oxford
[2]Physiology, Anatomy and Genetics
[3]Max-Planck Institute for Dynamics and Self-Organization

Decades of sensory neuroscience have begun to unravel how information about the outside world is encoded in sensory areas in order to inform subsequent action. Neural information encoding is often formulated in terms of tuning curves, pairwise correlations between neurons, or projection of neural activity onto a lower dimensional

space. However, it is unclear which features of neural activity enable robust propagation of activity, which is crucial to understand given that activity needs to be transmitted amongst brain regions to guide behaviour. To address this, we imaged two hierarchically organised and densely interconnected regions, the primary and secondary somatosensory cortex (S1 and S2) while performing two-photon photostimulation of S1 neurons. We assigned behavioural salience to photostimulation by training mice through operant conditioning to report detection of photostimulation. Photostimulation was targeted to groups of 5-50 neurons in S1, allowing us to parametrically vary the effect of photostimulation on behaviour. Hence we were able to address which features of activity in S1 explain propagation of activity to S2 and/or behaviour. We demonstrate that behaviourally salient photostimulation of S1 (hit trials) elicits robust amplified activation of S2, hinting that amplification may be a strategy employed by cortex to propagate sparse activity. Further, we analysed the covariance structure of S1 activity immediately preceding stimulation. The covariance structure of sensory populations is thought to relate to attention (Huang et al. Neuron 2019) and has been hypothesised to underpin activity propagation in silico (Zylberberg et al. PLoS Comp. Bio. 2017). Strikingly, we find that high dimensionality of population activity prior to the stimulus is predictive both of activity propagation from S1 to S2 and behavioural outcome. Hence our work demonstrates that the dimensionality of population activity may be critical to sensory processing as it both facilitates information propagation downstream and informs subsequent behaviour.

T-26. Optimization for real-world environments explains biological sound localization

Andrew Francl[1,2] FRANCL@MIT.EDU
Josh McDermott[1] JHM@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Department of Brain and Cognitive Sciences

Mammals localize sounds using information from their two ears. Localization in real-world conditions is challenging, as echoes provide erroneous information, and noises mask parts of target sounds. To better understand how real-world localization might be adapted to these challenges, we equipped a deep neural network with human ears and trained it to localize sounds in a virtual environment. The resulting model localized accurately in realistic conditions with noise and reverberation. In simulated experiments, the network exhibited many features of spatial hearing: sensitivity to monaural spectral cues and interaural time and level differences, integration across frequency, and biases for sound onsets. But when trained in unnatural environments without either reverberation, noise, or natural sounds, these performance characteristics deviated from those of humans. The results show how biological hearing is adapted to the challenges of real-world environments and illustrate how artificial neural networks can extend traditional ideal observer models to real-world domains.

T-27. Sinusoidal transformation of the visual field on the cortical surface

Madineh Sedigh-Sarvestani MADINEH.SARVESTANI@MPFI.ORG
David Fitzpatrick DAVID.FITZPATRICK@MPFI.ORG

Max Planck Florida Institute of Neuroscience

Mammalian cortex contains multiple topographic maps of the sensory environment wherein nearby neurons represent nearby locations in the environment locally and globally. In the visual system, the surface of the retina is transformed onto the surface of cortical areas, resulting in a 'retinotopic map'. Our study reveals a fundamentally new class of such maps that questions long-held assumptions about the principles of cortical sensory surface representation. We show retinotopic maps in the elongated secondary visual area of the tree shrew that exhibit a complex sinusoidal relationship to the retinal surface preserving only local relationships, an organization incompatible with simple mapping prevalent in the published literature. We use mathematical models designed to achieve energy efficiency to demonstrate that sinusoidal maps are the optimal solution in elongated areas. Furthermore, we show that this sinusoidal topography is accompanied by oscillating patterns of intracortical connections and stimulus response properties. Our findings suggest that cortical wiring flexibly implements solutions to energy minimization, with drastic influences on topographic maps and neuronal properties.

T-28. The geometry of the representation of decision variable and stimulus difficulty in the parietal cortex

Gouki Okazawa[1,2] OKAZAWA@NYU.EDU
Christina Hatch[1] CEH440@NYU.EDU
Allan Mancoo[3] ALLAN.MANCOO@NEURO.FCHAMPALIMAUD.ORG
Christian Machens[3] CHRISTIAN.MACHENS@NEURO.FCHAMPALIMAUD.ORG
Roozbeh Kiani[1] ROOZBEH@NYU.EDU

[1]New York University
[2]Center for Neural Science
[3]Champalimaud Centre for the Unknown

Lateral intraparietal (LIP) neurons represent formation of perceptual decisions involving eye movements. In circuit models proposed for these decisions, neural ensembles that encode actions compete to form decisions. Consequently, decision variables (DVs) are represented as partially potentiated action plans, where ensembles increase their average responses for stronger evidence supporting their preferred actions, as has been observed in LIP during motion direction discrimination tasks. As another consequence, DV representation and readout are implemented similarly for decisions with identical competing actions, irrespective of input and task context differences. Here, we challenge those core principles using a novel face discrimination task. We found that LIP firing rates decrease with supporting evidence, contrary to conventional motion discrimination tasks. These opposite response patterns arise from similar mechanisms in which decisions form along curved population-response manifolds. The curved manifolds rotate in state space depending on task context, necessitating task-dependent readouts. The encoding of the DV is misaligned with the action preferences of neurons, inconsistent with assumptions of past circuit models. We show similar manifolds in lateral and medial prefrontal cortices, suggesting a ubiquitous representational geometry across decision-making circuits.

T-29. Flexible identification of neural population dynamics underlying decision making

Mikhail Genkin[1] MGENKIN@CSHL.EDU
Chandramouli Chandrasekaran[2] CCHANDR1@BU.EDU
Krishna Shenoy[2] SHENOY@STANFORD.EDU
Tatiana Engel[1] ENGEL@CSHL.EDU

[1]Cold Spring Harbor Laboratory
[2]Boston University

Behaviorally relevant signals are often represented in neural population dynamics, which evolve on a low-dimensional manifold embedded into a high-dimensional space of neural responses. Revealing population dynamics from spikes is challenging because the dynamics and embedding are nonlinear and obscured by diverse responses of individual neurons and spiking noise. For example, the decision-related activity of single neurons was hypothesized to arise from either gradual ramping or abrupt stepping dynamics on single trials, but selection between these alternatives is brittle due to diversity of neural responses. Moreover, ramping and stepping are impoverished hypotheses for heterogeneous decision-related neural populations. We need frameworks that can flexibly identify neural dynamics from data. We developed a flexible framework for inferring neural population dynamics from spikes. In our framework, population dynamics are modeled by a latent dynamical system defined by a potential function. The framework covers a continuous space of hypotheses represented by different potential shapes. The activity of individual neurons is related to the population dynamics through unique firing-rate functions, which account for the heterogeneity of neural responses. We incorporate the distribution of the population state at the trial start and boundary conditions modeling trial termination, such as commitment to a choice when reaching a decision boundary. The potential, firing-rate functions, and initial state distribution are continuous functions that are inferred from data. On simulated neurons, our framework correctly recovered the ramping and stepping models, which correspond to linear and three-well potentials, respectively. We applied the framework to neurons recorded simultaneously from the macaque premotor cortex during decision making. The inferred potential revealed dynamics that evolve gradually towards the correct choice but have to overcome a potential barrier towards the incorrect choice, inconsistent with the simple hypotheses proposed previously. Our results demonstrate that a flexible approach can discover new hypotheses about population dynamics from data.

T-30. Attention improves information flow between neuronal populations without changing the communication subspace

Ramanujan Srinath[1,2] RAMANUJAN@PITT.EDU
Douglas A Ruff[1] RUFFD@PITT.EDU
Marlene R Cohen[1] COHENM@PITT.EDU

[1]University of Pittsburgh
[2]Center for Neural Basis of Cognition

Perhaps the most impressive hallmark of the nervous system is its flexibility. We effortlessly alternate between relying on or ignoring the same sensory information in different contexts. In the visual system, a population of neurons sends projections to a variety of different sensory, association, and motor areas, and functional studies show that only a small proportion of neuronal population activity is shared between even highly connected brain areas. The goal of our study was to test the hypothesis that this divergent functional (or anatomical) connectivity is a substrate for flexible, attention-dependent routing of sensory information. We recorded simultaneously from dozens of visual neurons in the medial temporal area (MT) and oculomotor neurons in the superior colliculus (SC) with overlapping receptive fields while rhesus monkeys performed a task in which they switched spatial attention, alternatingly using or ignoring the stimulus in the joint receptive fields of the recorded neurons. We investigated three potential mechanisms of flexible information flow, testing the hypotheses that attention modulates information flow between areas by (1) changing local representations, (2) changing the communication subspace, and/or (3) changing the efficacy of information transfer. We found that trial-to-trial variability of the population of MT neurons was better predicted by the activity of SC neurons (and vice versa) when attention was directed inside their joint receptive fields. Surprisingly, this improvement in prediction was not accompanied by changes in the dimensionality of the shared subspace and was not explained by increases or decreases in private or shared pairwise noise correlations. These results are the first demonstration of how attention affects the functional communication of visual information between populations of neurons across visual areas and suggests a mechanism as to how cognitive processes can affect perceptual decision making in ways that are independent of changes to the local neuronal representations.

1-001. Spontaneous and evoked activity patterns diverge over development

Lilach Avitan[1,2] LILACH.AVITAN@MAIL.HUJI.AC.IL
Zac Pujic[1] Z.PUJIC@UQ.EDU.AU
Jan Molter[1] JAN.MOELTER@UQ.NET.AU
Shuyu Zhu[1] I.ZHU@UQ.EDU.AU
Biao Sun[1] B.SUN1@UQ.EDU.AU
Geoffrey Goodhill[1] G.GOODHILL@UQ.EDU.AU

[1]The University of Queensland
[2]Queensland Brain Institute

The immature brain is highly spontaneously active. Over development this activity must be integrated with emerging patterns of stimulus-evoked activity, but little is known about how this occurs. Here we investigated this question by recording spontaneous and evoked neural activity in the larval zebrafish tectum from 4 to 15 days post fertilisation. Correlations within spontaneous and evoked activity epochs were comparable over development, and their neural assemblies refined in similar ways. However both the similarity between evoked and spontaneous assemblies, and also the geometric distance between spontaneous and evoked patterns, decreased over development. At all stages of development evoked activity was of higher dimension than spontaneous activity. Thus spontaneous and evoked activity do not converge over development in this system, and these results do not support the hypothesis that spontaneous activity evolves to form a Bayesian prior for evoked activity.

1-002. Dynamics of a double ring model of motor cortex during movement preparation and execution

Ludovica Bachschmid Romano[1,2] LUDOVICA.BACHSCHMID.ROMANO@DUKE.EDU
Nicholas G Hatsopoulos[3] NICHO@UCHICAGO.EDU
Nicolas Brunel[4] NICOLAS.BRUNEL@DUKE.EDU

[1]Duke University School of Medicine
[2]Neurobiology
[3]University of Chicago

[4]Duke University

The role of the activity of motor cortex during the stages of motor planning that precede the execution of movement remains a central question that must be answered to understand the neural substrates of motor control. The results of delayed-reach task experiments have shown that neurons in the primary motor cortex (M1) respond selectively to movement directions during the delay period, despite no movement occur- ring. During movement preparation and execution, neural activity has been shown to be confined to subspaces that are approximately mutually orthogonal. The network mechanisms that underlie these properties re- main poorly understood. We developed a firing-rate network model that is able to reproduce the temporal evolution of multi-unit activity recorded from M1 of a macaque monkey during an instructed delayed-reach task. The model assumes that the selectivity of a single neuron response to the direction of movement dy- namically changes throughout the course of the motor action, even though the direction of movement that is encoded at the population level remains stable. Our analysis suggests that during the delay period, information encoding the direction of movement is inherited by the network through external inputs that are tuned to the preferred directions of the neurons; during movement execution, the activity is instead driven by an untuned external input, and information regarding the movement direction is maintained via strong direction-specific recurrent connections. Simulations of network dynamics show that the degree of orthogonality of the prepara- tory and execution subspaces are in good agreement with the data. Altogether, our model explains how activity during the preparatory stage can set the direction of movement that will be encoded by the network throughout the entire duration of the motor action while simultaneously avoiding movement-potent dimensions, providing new insights into the relationship between motor planning and motor execution in M1.

1-003. Neural sampling from bimodal distributions in primary visual cortex

Ruben Coen-Cagli RUBEN.COEN-CAGLI@EINSTEINMED.ORG
Adam Kohn ADAM.KOHN@EINSTEINMED.ORG

Albert Einstein College Of Medicine

The neural sampling theory proposes that fluctuations of neural activity across multiple presentations of a stimulus (response variability) reflect the uncertainty of probabilistic inferences. The speed with which neurons can draw independent samples determines the accuracy of sampling-based neural codes. Widespread dynamical features of cortical activity, including stimulus-dependent oscillations and transients, appear consistent with fast sampling algorithms [1]. Here we consider a key test case, namely the representation of bimodal probability distributions. We test whether neurons in macaque primary visual cortex (V1) sample from a single mode during each individual stimulus presentation (slow sampling), or visit both modes within a single presentation (fast sampling). First, we identified a class of images that elicits bimodal probability distributions. Recent work shows that spatial context (stimuli outside the neural receptive field, RF) suppresses response variability, reflecting that contextual information reduces uncertainty [2]. We find that response variability is weaker when spatial context is informative about the RF stimulus (statistically homogeneous; [3]) than when it is not (heterogeneous). However, images with equal probability of being homogeneous or heterogeneous (termed ambiguous) evoke higher variability than either class. We show that neural sampling from bimodal distributions explains this increased variability. We then studied the within-trial dynamics of responses to ambiguous images. We find that in most cases (N=135/168 neurons and images) the data are significantly better explained by a model that switches between modes within each trial (Hidden Markov Model with Poisson observations), than by a model that samples around a single mode within each trial. Our results offer a strong test of neural sampling, showing for the first time that V1 response statistics may represent more than just first and second order moments of the posterior distribution, and supporting the view that V1 dynamics is consistent with fast sampling algorithms.

1-004. Understanding the role of sparseness in cerebellar granule cell representations

Marjorie Xie[1,2] MARJORIE.XIE@COLUMBIA.EDU
Samuel Muscinelli[1,3] SPM2176@COLUMBIA.EDU
Kameron Decker Harris[4] KAMERON.HARRIS@WWU.EDU
Ashok Litwin-Kumar[1,2] AK3625@COLUMBIA.EDU

[1]Columbia University
[2]Neuroscience
[3]Zuckerman Institute
[4]Western Washington University

Cerebellar anatomy is remarkably uniform. Sensorimotor inputs from mossy fibers expand into a vast layer of granule cells (GCs) before converging onto Purkinje cells. This feedforward expansion motif repeats throughout the cerebellar cortex, suggesting a general-purpose computation. According to Marr-Albus theories, the primary function of GC populations is to separate patterns: two distinct stimuli can be decorrelated through sparse and high-dimensional representations before being linearly separated. However, these theories assume the tasks being learned are random mappings between inputs and outputs. In its biological context, the cerebellum is critical for executing smooth movements, often described as a forward model predicting movement dynamics. Understanding the cerebellum's learning capabilities thus requires a perspective beyond that of a random-pattern classifier. We therefore ask, what role do GC representations play in learning smooth input-output mappings? We investigate the dependence of generalization performance on circuit parameters such as the coding level of GC representations (fraction of neurons active). We find that optimal coding levels depend on task structure: smoother tasks favor models with higher coding levels, a departure from previous theories. To explain these observations, we leverage kernel theory methods. We find that higher coding levels form a representation better suited to learn target functions rich in low frequencies. These trends hold under stronger biological constraints on our model, including sparse in-degree connectivity of GCs. Our theory predicts a dependence of GC coding level on the frequency content of the task and offers a normative explanation for task-to-task differences in GC population coding levels observed in vivo. We propose that while anatomically homogeneous, the cerebellar cortex can be viewed as comprising flexible circuits whose inductive bias for task structures can be tuned through circuit parameters such as coding level.

1-005. Creative memory semantization through adversarial dreaming

Nicolas Deperrois[1,2] NICOLAS.DEPERROIS@UNIBE.CH
Jakob Jordan[3,4] JAKOB.JORDAN@UNIBE.CH
Mihai Petrovici[5,6] MIHAI.PETROVICI@UNIBE.CH
Walter Senn[1] WALTER.SENN@UNIBE.CH

[1]Universitat Bern
[2]Institut fur Physiologie
[3]University of Bern
[4]Department of Physiology
[5]University of Bern and Heidelberg University
[6]Department of Physiology and Kirchhoff-Institute for Physics

Classical theories of memory consolidation emphasize the importance of sleep-mediated replay as a key mechanism to extract semantic information from episodic memories. However, the creative nature of dreams suggest that efficient memory semantization may go beyond merely replaying previous experiences. Here, we suggest that rapid-eye-movement (REM) sleep reorganizes memory by randomly combining episodic memories and creating entirely new visual experiences. Non-REM (NREM) sleep, in contrast, is responsible for making internal representations robust via the perturbed replay of encoded memories. We support this hypothesis by implementing a cortical architecture with separate, hierarchically organized forward and backward pathways, loosely inspired by generative adversarial networks (GANs). During wakefulness, episodic memories are stored in hippocampus. During REM sleep, these memories are randomly combined in high-level areas to generate new, hypothetical activity patterns in low-level areas. While feedforward pathways learn to distinguish these internal dreams from externally driven inputs, feedback pathways adversarially learn to generate more realistic activity patterns. During NREM sleep, episodic memories generate perturbed activity in early sensory areas. Reconstructing the encoded memory from this activity improves the robustness of semantic memory to potential sensory perturbations. Our cortical architecture, trained on standard benchmark datasets, develops rich latent representations in an unsupervised fashion. Using these, a linear classifier achieves competitive recognition performance, on par with unsupervised machine learning methods. By systematically evaluating the quality of the learned representations, our results demonstrate the complementary function of NREM and REM sleep. The suggested architecture and learning paradigm are amenable to cortical microcircuit implementation in terms of different classes of layer 2/3 pyramidal neurons (Sacramento et al., 2018). Our biologically-inspired framework highlights the importance of creative dreaming for the successful semantization of memories, providing a new, functional perspective on the link between sleep, dreaming, hippocampal replay and memory consolidation.

1-006. Revisiting and revising clustering in brain networks

Tanguy Fardet[1] TANGUY.FARDET@TUEBINGEN.MPG.DE
Anna Levina[2] ANNA.LEVINA@UNI-TUEBINGEN.DE
Lukas Paulun[3] MAIL@LUKASPAULUN.DE

[1]University of Tubingen, Max Planck Institute for Biological Cybernetics
[2]Universtity of Tubingen
[3]Bernstein Center Freiburg, University of Freiburg

Nodes in real-world networks tend to cluster into densely connected groups, a property captured by the clustering coefficient. This non-random tendency for clustering is notably strong in neuronal networks, where connections are typically directed and were shown to have broad (e.g. lognormal) weight distributions. However, clustering was initially defined for unweighted and undirected networks, which leaves out crucial information about dynamics on weighted graphs. Several generalizations have been proposed for weighted networks but none of them fulfills the continuity condition: graph measures should not be influenced by the addition or deletion of edges having very small weights. This condition means that an edge with infinitesimally small weight should be equivalent to the absence of that edge. We propose here a new definition of the clustering coefficient that tackles this issue while satisfying previously formulated requirements. This new definition is less sensitive to weak spurious connections that are prevalent in inferred brain networks due to noise or statistical biases in measured graphs. Compared to previous methods, it is able to detect topological features more precisely and behaves in a more sensible manner for inferred networks. Finally, we discuss the differences between clustering methods for various brain networks, from C. elegans to the mouse brain. We discuss how our purely weighted continuous definition can differ from both the binary and other weighted definitions, and show how our analysis challenges several preconceptions regarding generic properties of brain networks.

1-007. Explaining dopaminergic response heterogeneity via high-dimensional value function approximation

Rachel Lee[1] RACHELSTEPHLEE@GMAIL.COM
Ilana Witten[2] IWITTEN@PRINCETON.EDU
Nathaniel Daw[2] NDAW@PRINCETON.EDU
Ben Engelhard[1] BENENGELHARD@PRINCETON.EDU

[1]Princeton Neuroscience Institute, Princeton University
[2]Princeton University

Phasic responses from midbrain dopamine (DA) neurons have been argued to report a reward prediction error (RPE) signal from reinforcement learning (RL) models. In the classic formulation, RPE is defined as a scalar signal, consistent with early work suggesting homogeneous responses across neurons. However, recent work has challenged this view by clearly demonstrating that DA can exhibit heterogeneous responses across neurons and regions. We examine one dataset reporting cell body responses during a complex visual evidence accumulation task. Two key features emerge: neurons respond heterogeneously to a variety of features (some seemingly RPE-irrelevant) during the cue period (ie the navigation and decision making portion of the task), but the population behaves largely uniformly with RPE-like responses at trial outcome. We introduce a new model that explains both aspects of the data by positing that DA neurons report individual RPEs for each dimension of a population vector code for the task's state. Given any such representation, this scheme reproduces both cue period heterogeneity (driven by the feature-dependent future value term in the RPE) and outcome period homogeneity (driven by a scalar reward). To investigate this claim, we train a deep RL model on the evidence accumulation task to extract a feature representation of task state from raw video input. The vector of RPEs derived from the features of the deep RL model exhibit similar qualitative behavior as the DA neurons: specifically, the vector RPEs respond to various sensory variables of the task during cue period but are still uniformly modulated by RPE at reward time. Taken together, our work provides a path to reconcile new observations of DA neuron heterogeneity with classic ideas about RPE coding, while also providing a new perspective on how the brain performs reinforcement learning in high dimensional environments.

1-008. Metacognition guides near-optimal exploration of a large state space with sparse rewards

Su Jin An SUJINAN@KAIST.AC.KR
Sang Wan Lee SANGWAN@KAIST.AC.KR

Korea Advanced Institute of Science and Technology

Earlier studies have used the reinforcement learning theory to explain how animals explore a task space to maximize reward. However, a majority of empirical tests heavily rely on a simple task paradigm. This limits our understanding of the ability to explore an uncharted world with infinitely-many options, which inevitably entails the sparse reward problem. Here, we test a theoretical idea that metacognition1,2, the ability to introspect and estimate one's own level of uncertainty, guides efficient exploration. We designed a novel two-stage decision-making task with infinitely-many choices and sparse rewards (90% rewarded in less than 8% options in reward states on average) and collected 88 subjects' behavioral data. First, we identified two key variables guiding exploration: uncertainty about the environmental structure (state-space uncertainty) and information about the reward structure (reward information). To further understand exploration dynamics, we differentiated between the two variables as a function of a learning stage. We found that the state-space uncertainty is significantly correlated with the individual metacognitive ability measured using an independent perception task3. Interestingly, highly metacognitive subjects act on the state-space uncertainty over the course of learning, while the effect of the reward information on exploration behavior diminishes after the early learning stage. Note the learning bias towards the environmental structure and against the reward structure is a near-optimal exploration strategy for the sparse reward problem. On the other hand, the effects of both variables last in the low metacognitive subject group. Our theory is further supported by the finding that the high metacognitive subject group showed higher task performance and sampling efficiency in the test phase following the learning phase. Taken together, our work elucidates a crucial role of metacognition in fostering a sample-efficient, near-optimal exploration strategy to resolve uncertainty about environmental and reward structures.

1-009. Deep generative analysis for task-based functional MRI experiments

Daniela de Albuquerque[1,2] DFD4@DUKE.EDU
Jack Goffinet[1] JACK.GOFFINET@DUKE.EDU
Rachael Wright[1] RACHAEL.WRIGHT@DUKE.EDU
John Pearson[1,3] JOHN.PEARSON@DUKE.EDU

[1]Duke University
[2]Department of Electrical and Computer Engineering
[3]Biostatistics & Bioinformatics

While functional magnetic resonance imaging (fMRI) remains one of the most widespread and important methods in basic and clinical neuroscience, the data it produces—time series of brain volumes—continue to pose daunting analysis challenges. The current standard ("mass univariate") approach involves constructing a matrix of task regressors, fitting a separate general linear model at each volume pixel ("voxel"), computing test statistics for each model, and correcting for false positives post-hoc using bootstrap or other resampling methods. Despite its simplicity, this approach has enjoyed great success over the last two decades due to:1) its ability to produce effect maps highlighting brain regions whose activity significantly correlates with a given variable of interest; and 2) its modeling of experimental effects as separable and thus easily interpretable. However, this approach suffers from several well-known drawbacks, namely: inaccurate assumptions of linearity and noise Gaussianity; a limited ability to capture individual effects and variability; and difficulties in performing proper statistical testing secondary to independently fitting voxels. In this work, we adopt a different approach, modeling entire volumes directly in a manner that increases model flexibility while preserving interpretability. Specifically, we use a generalized additive model (GAM) in which the effects of each regressor remain separable, the product of a spatial map produced by a variational autoencoder and a (potentially nonlinear) gain modeled by a covariate-specific Gaussian Process. The result is a model that yields group-level effect maps comparable or superior to the ones obtained with standard fMRI analysis software while also producing single-subject effect maps capturing individual differences. This suggests that generative models with a decomposable structure might offer a more flexible alternative for the analysis of task-based fMRI data.

1-010. Non-linear temporal dynamics of neuronal responses in human visual cortex

Iris Groen[1,2] I.I.A.GROEN@UVA.NL
Giovanni Piantoni[3] G.PIANTONI@UMCUTRECHT.NL
Stephanie Montenegro[4] STEPHANIE.MONTENEGRO@NYULANGONE.ORG
Adeen Flinker[4] ADEEN.FLINKER@NYULANGONE.ORG
Sasha Devore[4] SASHA.DEVORE@NYULANGONE.ORG
Orrin Devinsky[4] ORRIN.DEVINSKY@NYULANGONE.ORG
Werner Doyle[4] WERNER.DOYLE@NYUMC.ORG
Nick Ramsey[3] N.F.RAMSEY@UMCUTRECHT.NL
Natalia Petridou[3] N.PETRIDOU@UMCUTRECHT.NL
Jonathan Winawer[5] JONATHAN.WINAWER@NYU.EDU

[1]University of Amsterdam
[2]Institute for Informatics
[3]University Medical Center Utrecht
[4]New York University School of Medicine
[5]New York University

Neural responses in visual cortex exhibit complex temporal dynamics. For example, responses decrease when a static visual stimulus is prolonged in time (subadditive temporal summation), reduce to stimuli that are repeated (adaptation), and rise less rapidly for low contrast stimuli (slow dynamics). These dynamics are typically measured and modeled with a tailored set of stimuli designed to investigate one particular kind of temporal phenomenon (e.g. adaptation), often using different measurement techniques and different computational models. We present a new dataset of visual neural response measurements in which multiple temporal phenomena were measured simultaneously by systematically varying three different dimensions: duration (to measure temporal summation), inter-stimulus-interval (to measure neural adaptation) and contrast (to measure temporal slowing). The data are electrocorticographic (ECoG) recordings of human visual cortex, which track neural responses at the millisecond scale with high spatial and temporal precision. By mapping electrodes to probabilistic retinotopic atlases within individual participants and aggregating measurements over multiple participants, we collected a large and comprehensive set of responses across multiple visual areas, covering both earlier (V1-V3a/b) and higher-order (LO, TO, IPS) retinotopic maps. We show that a single computational model consisting of a small set of canonical neural computations (linear filtering, rectification, exponentiation, and delayed normalization) predicts the observed response variations across all three stimulus manipulations with high accuracy. Systematic removal of each of these canonical computations from the full model revealed that non-linear computations, and in particular the delayed normalization, are critical to simultaneously account for temporal summation, adaptation, and slower dynamics at low contrast. Furthermore, comparison of model predictions across the visual hierarchy revealed increased temporal window lengths and more pronounced nonlinearities in higher visual regions. These results demonstrate that non-linear canonical neural computations capture a wide range of temporal and contrast-dependent neuronal dynamics at millisecond resolution throughout visual cortex.

1-011. Learning from unexpected events in the neocortical microcircuit

Jason E. Pina[1,2] JAYPINA@YORKU.CA
Colleen J Gillon[3] COLLEEN.GILLON@MAIL.UTORONTO.CA
Jerome Lecoq[4] JEROMEL@ALLENINSTITUTE.ORG
Jed Perkins[4] JEDP@ALLENINSTITUTE.ORG
Sam Seid[4] SAMS@ALLENINSTITUTE.ORG
Carol Thompson[4] CAROLTH@ALLENINSTITUTE.ORG
Ryan Valenza[4] RYANV@ALLENINSTITUTE.ORG
Tim M Henley[1,5] HENLEYT@YORKU.CA
Yoshua Bengio[6] YOSHUA.BENGIO@MILA.QUEBEC
Timothy P Lillicrap[7] TIMOTHYLILLICRAP@GOOGLE.COM
Joel Zylberberg[1,8] JOEL.ZYLBERBERG@UCDENVER.EDU
Blake A Richards[9] BLAKE.RICHARDS@MCGILL.CA

[1]York University
[2]Physics and Astronomy
[3]University of Toronto
[4]Allen Institute for Brain Sciences
[5]Physics
[6]Montreal Institute for Learning Algorithms
[7]Google DeepMind
[8]Centre for Vision Research

[9]McGill University

Whether neocortex learns the structure of a changing environment in a predictive, hierarchical manner is a long-standing question in neuroscience. To do so, expected, predictable features must be differentiated from unexpected, unpredictable ones by comparing bottom-up and top-down data streams. The cortex must then change the representation of these stimuli in response to mismatch signals that arise from the unexpected features. While such mismatch signals and differences in bottom-up versus top- down information have been partially investigated in prior work, to our knowledge whether the differences govern the subsequent evolution of feature representations and, thus, of learning remains unexplored. Here, we present results demonstrating that mismatch signals can determine changes in responses to expected and unexpected stimuli in individual neuronal responses that are tracked over a period of days. We imaged layer 2/3 (L2/3) and L5 pyramidal neurons in primary visual cortex in awake mice using 2-photon calcium imaging. To explore the differences in bottom-up and top-down signals, we recorded the activity of somata and distal apical dendrites, which are electrotonically separated from the somata and receive top-down feedback. We examined how these signals varied with unexpected events by having the animals passively view “expected” sequences of frames composed of Gabor patches that were substituted with “unexpected” sequences at random times during each session. We found (1) many neurons in both layers showed large differences in their responses to expected and unexpected events, (2) these differences evolved over days, (3) mismatches in responses determined how the differences evolved, and (4) the way the responses evolved varied between distal apical dendrites and somata, as expected for hierarchical models. Together, our results suggest that visual cortex instantiates a predictive hierarchical model in which unexpected events drive learning.

1-012. Distributed local learning in Context-Switched Linear Networks

SueYeon Chung SUEYEON.CHUNG@COLUMBIA.EDU
Larry Abbott LFA2103@COLUMBIA.EDU

Columbia University

Recent experimental findings have stressed that the signals used for controlling neural plasticity can be different from those propagated to downstream targets (Muller et al, 2019). Inspired by these findings, we propose Context-Switched Linear Networks (CSLNs), a neural network architecture that employs context-dependence via sparsity and has separate representations for learning and inference. This architecture has a close connection to Gated Linear Networks (GLNs) (Veness et al, 2017), a recently proposed family of deep networks that gain computational power from context-switching rather than from nonlinear transfer functions. We show that this family of architectures can: 1) employ ‘expert’ neurons that performs inference for a categorization task directly, but in a distributed manner, 2) learn locally via online convex optimization, and 3) show superior performance in tasks that are known to be difficult for backpropagation, such as multi-task learning.

1-013. Ecological pre-training of RNNs recovers rats suboptimal behavior in a simple categorization task

Manuel Molano-Mazon[1,2] MANUELMOLANOMAZON@GMAIL.COM
Daniel Duque[3] DDDBCN@GMAIL.COM
Guangyu Robert Yang[4] GYYANG.NEURO@GMAIL.COM
Jaime de la Rocha[1] JROCHAV@CLINIC.CAT

[1]IDIBAPS
[2]Brain Circuits and Behavior lab
[3]Institut d’Investigacions Biomediques August Pi i Sunyer (IDIBAPS)
[4]Columbia University

In order to make optimal decisions in real-world environments, it is fundamental to take into account previous experiences. A recent study (Hermoso-Mendizabal et al. 2020) has shown that rats leverage the information contained in the trial history of an auditory two-alternative forced choice (2AFC) task that includes trial-to-trial correlations. However, rats failed to use this trial history information after errors, when they based their choices only on the current stimulus. Here we first confirm that this reset strategy consistently arises in several task variants. We then train Recurrent Neural Networks (RNNs) in a similar 2AFC task and show that, like the rats, the networks can leverage the sequence partial predictability to increase their categorization accuracy after correct trials. However, RNNs can also capitalize on this predictability after error trials, departing from the reset strategy of the rats, and resembling the behavior of an optimal agent that can perform counterfactual inference and reverse its prediction about the rewarded port after errors. We propose that this discrepancy may reflect the adaptation of the rat’s brain to more complex, ecological environments in which, due to the large number of alternatives, there is a fundamental asymmetry in the information provided by correct and error choices. As a consequence, rats disregard the predictive power of unrewarded choices in the 2AFC task. We show that, if RNNs are pre-trained in a task using a larger number of choices N > 2, they exhibit the same reset strategy displayed by the rats when tested in the 2AFC task. Our results suggest that seemingly sub-optimal behaviors in laboratory tasks (e.g. the reset strategy) can be a byproduct of the adaptation to more complex environments. Furthermore, our work demonstrates the importance of pre-training artificial neural networks in ecologically plausible environments before testing them on the task of interest.

1-014. Deep Graph Pose: a semi-supervised deep graphical model for improved animal pose tracking

Anqi Wu[1] aw3236@columbia.edu
E. Kelly Buchanan[1,2] ekb2154@columbia.edu
Matthew R. Whiteway[1] mw3323@columbia.edu
Michael Schartner[3] michael.schartner@unige.ch
Guido Meijer[4] guido.meijer@research.fchampalimaud.org
Jean-Paul Noel[5] jpn5@nyu.edu
Erica Rodriguez[1] er2934@columbia.edu
Claire Everett[1] cpe2108@columbia.edu
Amy Norovich[1] aln2128@columbia.edu
Evan Schaffer[1] ess2129@columbia.edu
Neeli Mishra[1] nm2786@columbia.edu
C. Daniel Salzman[1] cds2005@columbia.edu
Dora Angelaki[5] da93@nyu.edu
Andres Bendesky[1] ab4463@columbia.edu
The International Brain Laboratory[6] info@internationalbrainlab.org
John Cunningham[1] jpc2181@columbia.edu
Liam Paninski[1] lmp2107@columbia.edu

[1]Columbia University
[2]Neuroscience
[3]University of Geneva
[4]The Champalimaud Centre for the Unknown
[5]New York University
[6]The International Brain Laboratory

Noninvasive behavioral tracking of animals is crucial for many scientific investigations. Recent transfer learning approaches for behavioral tracking have considerably advanced the state of the art. Typically these methods treat each video frame and each object to be tracked independently. In this work, we improve on these methods (particularly in the regime of few training labels) by leveraging the rich spatiotemporal structures pervasive in behavioral video — specifically, the spatial statistics imposed by physical constraints (e.g., paw to elbow distance), and the temporal statistics imposed by smoothness from frame to frame. We propose a probabilistic graphical model built on top of deep neural networks, Deep Graph Pose (DGP), to leverage these useful spatial and temporal constraints, and develop an efficient structured variational approach to perform inference in this model. The resulting semi-supervised model exploits both labeled and unlabeled frames to achieve significantly more accurate and robust tracking while requiring users to label fewer training frames. In turn, these tracking improvements enhance performance on downstream applications, including robust unsupervised segmentation of behavioral “syllables,” and estimation of interpretable “disentangled” low-dimensional representations of the full behavioral video.

1-015. Invariant texture recognition in mouse visual cortex

Fengtong Du dufengtong@gmail.com
Sonia Joseph soniamollyjoseph@gmail.com
Marius Pachitariu pachitarium@janelia.hhmi.org
Carsen Stringer stringerc@janelia.hhmi.org

HHMI Janelia Research Campus

The discrimination of visual textures, irrespective of rotation, scale and viewpoint, can help animals perform a variety of visual tasks such as object recognition and object segmentation. The neural basis of such invariant texture recognition is largely unknown. We recorded ~40,000 neurons simultaneously in mouse visual cortex

while presenting ~14,000 visual stimuli sampled from 32 texture classes via random rotation, scaling and cropping. We found that the neural responses encoded the texture class robustly and a linear classifier trained on the neural data achieved test performance on single trials of 81.67% $\pm$ 2.74% (stddev, n=4 recordings, chance=0.03). Furthermore, we found that the patterns of errors in the task were highly-consistent across mice, and classification accuracy improved slightly in higher-order visual areas. These computations were supported by a subset of texture-coding neurons which formed 10% of the population and were spread throughout visual areas. Classification accuracy from the texture-coding neurons was similar to that from the entire population. To determine the computational operations leading to texture invariance, we next analyzed the responses of artificial neurons from a pretrained AlexNet model. These artificial neurons performed well in the texture classification task, but their pattern of errors did not match the patterns from the neural data. To better model the neural data, we next fit a CNN encoding model directly to the responses of the coding neurons. The model explained some of the response patterns in individual images, but it also did not reproduce the pattern of errors in the texture classification task. In summary, both artificial and neural populations encode visual textures in an invariant way, but the texture representations appear substantially different between real and artificial neurons.

1-016. A no-report paradigm reveals that face cells multiplex consciously perceived and suppressed stimuli

Janis Hesse[1,2] JANISHESSE@GOOGLEMAIL.COM
Doris Tsao[1] DORTSAO@CALTECH.EDU

[1]California Institute of Technology
[2]Biology and Biological Engineering

Having conscious experience is arguably the most important reason why it matters to us whether we are alive or dead. A powerful paradigm to identify neural correlates of consciousness is binocular rivalry, wherein a constant visual stimulus evokes a varying conscious percept. It has recently been suggested that activity modulations observed during rivalry may represent the act of report rather than the conscious percept itself. Here, we performed single-unit recordings from face patches in macaque inferotemporal (IT) cortex using a novel no-report paradigm in which the animal's conscious percept was inferred from eye movements. These experiments reveal two new results concerning the neural correlates of consciousness. First, we found that large proportions of IT neurons represented the conscious percept even without active report. Using high-channel recordings, including a new 128-channel Neuropixels-like probe, we were able to decode the conscious percept on single trials. Second, we found that even on single trials, modulation to rivalrous stimuli was weaker than that to unambiguous stimuli, suggesting that cells may encode not only the conscious percept but also the suppressed stimulus. To test this hypothesis, we varied the identity of the suppressed stimulus during binocular rivalry. We found that we could decode not only the conscious percept but also the suppressed stimulus from neural activity. Moreover, the same cells that were strongly modulated by the conscious percept also tended to be strongly modulated by the suppressed stimulus. Together, our findings indicate that (1) IT cortex possesses a true neural correlate of consciousness even in the absence of report, and (2) this correlate consists of a population code wherein single cells multiplex representation of the conscious percept and veridical physical stimulus, rather than a subset of cells perfectly reflecting consciousness. We propose two computational models for how this multiplexing may be implemented.

1-017. Non-linear computations in spiking low-rank neural networks

Ljubica Cimesa LJUBICA994@GMAIL.COM
Lazar Ciric LAZAR.CIRIC@ENS-CACHAN.FR
Srdjan Ostojic SRDJAN.OSTOJIC@ENS.FR

Ecole Normale Superieure

A central question in neuroscience is how networks of neurons perform computations. To address this question, an influential approach has been to use tools developed in the field of Artificial Intelligence, and train recurrent neural networks (RNNs) to perform a range of sensory, motor and cognitive tasks studied in systems neuroscience [Sussilo 2014]. This approach however faces two important challenges: trained RNNs are in general (i) difficult to interpret in terms of mechanisms; (ii) difficult to relate with data as they lack biological constraints. Concerning interpretability, a promising way forward are low- rank networks, a sub-class of RNNs that are mathematically tractable while being able to implement a vast range of tasks [Mastrogiuseppe and Ostojic 2018]. Here we examine how biological constraints can be added to low-rank RNNs by extending this class to spiking networks. We consider a randomly connected network of excitatory and inhibitory integrate-and-fire neurons [Brunel 2000], to which we add a low-rank connectivity structure. We first compare the dynamics in this network to a rate network with statistically identical low-rank structure. We show that external inputs elicit analogous low-dimensional dynamics in the two networks. We next turn to non-linear dynamics generated by positive feedback. We show that if the low-rank structure is zero-mean, the non-linear dynamics are independent of the details of the single-neuron transfer function, which is non-negative in spiking and symmetric around 0 in rate networks. Specifically, positive feedback leads to an identical bifurcation in spiking and hyperbolic tangent rate networks. We finally demonstrate that the resulting low-dimensional dynamics can be exploited to perform non-linear computations such as the detection of a stimulus in noise. Altogether, our results open the door to building interpretable spiking networks beyond random connectivity, that implement a large range of computational tasks.

1-018. Universality of interval variability constraints in the sequential activity of motor circuits

Alicia Garrido-Pena[1,2] ALICIA.GARRIDO@UAM.ES
Irene Elices[3] IRENE.ELICES@INSERM.FR
Rafael Levi[1] RLEVI66@GMAIL.COM
Francisco B. Rodriguez[1] F.RODRIGUEZ@UAM.ES
Pablo Varona[1,4] PABLO.VARONA@UAM.ES

[1]Universidad Autonoma de Madrid
[2]Departamento Ingenieria Informatica. Computational Neuroscience Group
[3]Institut de la Vision, Sorbonne Universite
[4]Ingenieria Informatica

Many neural circuits produce sequential activations that are directly related to behavior and are characterizable in terms of the intervals that build such sequences. This characterization goes beyond the traditional analysis of frequency and synchronization in neural rhythms, which contributes to unveil important aspects of transient neural dynamics underlying behavior. Central Pattern Generators (CPGs) are circuits that generate and coordinate motor movements by producing patterned sequences of activations in their constituent neurons. These robust rhythms are yet flexible and the time intervals that build the neural sequences adapt as a function of the behavioral context. The presence of robust dynamical invariants in the form of cycle-by-cycle linear relationships between two intervals of the pyloric CPG sequence and the period has been recently revealed. This work characterizes the variability of the intervals that build the rhythm and the associated sequence of the feeding CPG of the mollusk Lymnaea stagnalis. The study assesses both the activity obtained in long electrophysiological recordings of living neurons and the rhythm produced by a realistic conductance-based model. We first quantified the variability of all intervals building the sequence and then we characterized the relationships between specific intervals and the period in the form of dynamical invariants. To induce variability in the CPG model, we used current injection, following the stimulation used in experimental recordings. We report the presence of distinct constrained variability in the sequence time intervals and the existence of several dynamical invariants, which are modulated by the stimulation. The presence of robust dynamical invariants in CPG sequences, not only in the model but also in two animal species points out the universality of this phenomena.

1-019. Expected uncertainty drives phasic and tonic exploration during probabilistic learning

Sylvia Bindas[1,2] BINDAS@OHSU.EDU
Vincent D. Costa[1] COSTAV@OHSU.EDU

[1]Oregon Health and Science University
[2]Behavioral Neuroscience

In changing environments, adaptive decision-making requires balancing when to choose familiar, known options with when to explore new, unknown options. This balancing act, known as the explore-exploit tradeoff, is critical to how we make choices that can maximize reward. Specifically, exploration supports optimal decision-making by reducing the uncertainty associated with certain choices. Exploration is often seen as phasic, where the decision to explore depends on peaks in uncertainty that signal when the benefit of exploring is greatest. However, exploration can also be tonic, occurring more regularly in time. While tonic exploration has been demonstrated in settings where uncertainty is limited to discrete, unexpected rule changes, it is unclear how tonic exploration relates to expected uncertainty from stochastic reward outcomes. Here, we use a Bayesian modeling approach to show that spontaneous errors (i.e. lapses) in the reversal phase of a two-armed bandit reversal learning task reflect a form of tonic exploration. This tonic exploration coexists with phasic exploration in the task, and does not scale with environmental uncertainty. Further, we find that tonic exploration is directed rather than random, as lapses can be accurately predicted by Bayesian estimates of unpredictability and choice consistency.

Our results demonstrate how tonic exploration complements phasic exploration in changing environments as a directed strategy to reduce uncertainty and maximize reward.

1-020. Inferring latent dynamics underlying neural population activity via neural differential equations

Timothy Kim[1] TDKIM@PRINCETON.EDU
Thomas Luo[1] ZHIHAOL@PRINCETON.EDU
Jonathan Pillow[1] PILLOW@PRINCETON.EDU
Carlos Brody[2] BRODY@PRINCETON.EDU

[1]Princeton Neuroscience Institute, Princeton University
[2]Princeton Neuroscience Institute

An important problem in systems neuroscience is to identify the latent dynamics underlying neural population activity (Harvey et al., 2012; Churchland et al., 2012). One recent study showed that recurrent neural networks (RNN) can be trained to capture the nonlinear dynamics of population activity in a variety of brain areas, including motor cortex during reaching (Pandarinath et al., 2018). However, RNNs and their dynamics are generally high-dimensional, making it difficult to gain insight into the putatively low-dimensional dynamics underlying neural activity. Here we address this problem by introducing a low-dimensional nonlinear model for neural population dynamics based on neural ordinary differential equations (neural ODEs). Neural ODEs are a recent advance in scientific machine learning that use a deep neural network to directly model the time derivatives of a nonlinear system. The resulting derivatives are numerically integrated with robust ODE solvers (Chen et al., 2018). Using neural ODEs, we develop a framework to infer continuous low-dimensional latent flow fields and fixed points from spike trains. We apply our framework to a variety of synthetic datasets, and find that it accurately infers single-trial latent trajectories and flow fields. Unlike the RNN-based approach of (Pandarinath et al., 2018), which acausally infers the initial value of the latent dynamics from future spikes, our approach is able to infer initial values using spikes from previous trials. We apply our model to neural population activities in multiple areas within the medial frontal cortex of rats performing a pulse-based auditory decision-making task. Each pulse affects the dynamics of our model with a trainable discrete jump in the state (Jia and Benson, 2019), similar to the drift-diffusion model for pulsatile evidence (Author et al. [1]). Our model provides a general framework for investigating the neural mechanisms of decision-making and other cognitive computations through the lens of dynamical systems.

1-021. Comparative analyses of neural population dynamics across experimental subjects

Alex Williams[1,2] ALEX.H.WILLIA@GMAIL.COM
Erin Kunz[1] EKUNZ@STANFORD.EDU
Isabel Low[1,3] ILOW@STANFORD.EDU
Lisa Giocomo[1] GIOCOMO@STANFORD.EDU
Simon Kornblith[4] SKORNBLITH@GOOGLE.COM
Scott Linderman[1] SCOTT.LINDERMAN@STANFORD.EDU

[1]Stanford University
[2]Department of Statistics
[3]Neurobiology
[4]Google

Natural variation across individuals is a foundational concept of biology. Such variation is well-documented in motor movements and in biophysical circuit parameters. Yet, very little is understood about animal-to-animal variability at the intervening level of neural population dynamics in vertebrates. Low-dimensional visualizations of population activity often show intriguing geometric features, such as “rotations”, “tangling”, or attractor-like structures. Here we outline a systematic approach to compare these observations across neural recordings taken from different animal subjects or brain regions. To do this, we build upon ideas from statistical shape analysis to define notions of distance between population dynamics that are symmetric and obey the triangle inequality. Formally, such distance functions define a metric space, which enables a variety of cohort-level analyses. For example, given recordings from K subjects, we can characterize the extent to which their dynamics cluster or, alternatively, reflect a single “average.” In recurrent neural networks (RNNs) trained on a discrimination task, we show that averaging within the metric space reduces noise and results in interpretable low-dimensional trajectories. We also identify clusters across a population of RNNs with identical architectures. Next, we characterize position coding in mouse entorhinal cortex across “remapping” events. Across 18 Neuropixel recording sessions and 8 mice, we find that spatial maps extracted from the same session are most similar to each other. Strikingly, maps recorded from the same mouse on different days were often more similar than maps from other animals. Since each session was associated with a unique probe insertion, it is unlikely that this within-subject similarity is due to re-sampling the same neurons. Instead, this result may indicate reliable individual differences in circuit-level representations. Linking this variation in population dynamics to variation in behavioral outcomes (e.g. navigational strategies or performance) is one of several future directions enabled by our work.

1-022. Information-preserving dimensionality reduction for high-dimension, low-sample neural data

Aditi Jha[1,2] ADITIJH97@GMAIL.COM
Michael Morais[1] MJMORAIS@PRINCETON.EDU
Jonathan Pillow[3] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]Electrical Engineering
[3]Princeton Neuroscience Institute, Princeton University

The analysis of high-dimensional data is an important challenge in modern neuroscience. In many settings, it is desirable to reduce dimensionality while preserving information about some target variable of interest. For example, we might wish to identify the dimensions of a high-dimensional neural population response that carry information about a stimulus variable and discard all uninformative dimensions. Unsupervised methods like PCA cannot identify such dimensions because they do not take informativeness into account. This motivates sufficient dimension reduction (SDR) methods, which seek to identify a linear projection of high-dimensional data X that preserves the conditional distributions $X \mid Y$ for some target variable Y, discarding the dimensions of X that are statistically independent of Y [Globerson and Tishby, 2003]. However, existing SDR methods typically require more observations than the data-dimensionality ($N\>;p$), and their performance is highly sensitive to N. High-throughput neural recordings often fall into the opposite regime, where we have fewer trials than the number of neurons ($N\<;p$). This renders SDR methods intractable, even though evidence suggests that neural activity often varies meaningfully along only a small number of dimensions. Motivated by these considerations, we introduce Class-conditional Factor Analytic Dimensions (CFAD), a model-based dimensionality-reduction method for high-dimensional, small-sample-size datasets. The factor-analytic backbone of the model readily incorporates priors for known structure of the data, such as smoothness in the projection axes, which we show improve performance in the small-sample regime. We apply CFAD to simulated data and show that it dramatically outperforms existing methods in the low-sample, high-dimension regime. We then apply it to functional magnetic resonance imaging (fMRI) data and show that it outperforms all other methods including SIR, SAVE, LAD [Cook and Ni, 2005], Linear Discriminant Analysis, reduced-rank regression, and LOL [Vogelstein et al., 2017].

1-023. Quasi-Bayesian estimation of time constants supports lateralized auditory computation

Demetrios Neophytou[1,2] DNEOPHY000@CITYMAIL.CUNY.EDU
Diego Arribas[3,4] DIEGOARRI91@GMAIL.COM
Hysell Oviedo[1,2] HOVIEDO@CCNY.CUNY.EDU
Memming Park[3,4] MEMMING.PARK@STONYBROOK.EDU

[1]The City College of New York
[2]Biology
[3]Stony Brook University
[4]Neurobiology and Behavior

Recently, cortical asymmetry in auditory processing has been supported by behavioral and anatomical studies [1]. The left Auditory Cortex (ACx) is believed to specialize in processing fast temporal components of speech sounds, and the right ACx in slower components. To investigate possible circuit mechanisms underlying these asymmetries, we compared the connectivity of excitatory circuits in the left and right ACx. We found significantly stronger recurrent connectivity in the superficial layers of the right ACx compared to the left ACx. We hypothesize that the underlying recurrent neural dynamics would exhibit differential characteristic time scales corresponding to their hemispheric specialization. To investigate possible functional differences in temporal integration between the left and right ACx, we performed cell-attached recordings in awake, head-fixed mice, to compare the neural activity of superficial layers in response to frequency sweeps of varying speeds and directions. However, estimating the time constant from spike trains is challenging due to low firing rate neurons and the varying degrees of engagement with the network's time scale, resulting in large bias and large uncertainty in the estimate. Therefore, we developed a new quasi-Bayesian method to combine evidence from multiple weak signal-to-noise ratio neurons.

We use the dichotomized Gaussian model to reproduce autocorrelations and generate surrogate spike trains that were used to estimate uncertainty and provide a belief distribution over the time constant. Our method correctly infers the network time constant and provides (quasi-)credible interval even for a population of low signal-to-noise ratio neurons unlike the approaches used in the literature. Applying this method to the ACx data, we found a difference in the left and right ACx networks consistent with the time scale of hypothesized lateralized auditory signal processing.

1-024. Neural mechanisms of selection in visual working memory

Matthew Panichello[1,2] MFP2@PRINCETON.EDU
Timothy Buschman[1] TBUSCHMA@PRINCETON.EDU

[1]Princeton University
[2]Princeton Neuroscience Institute

Cognitive control is critical for flexible, adaptive behavior. For example, visual attention enhances the representation of goal-relevant stimuli, allowing this information to preferentially guide action. A similar mechanism is thought to control the selection of representations held ‘in mind’, in working memory (Myers, Stokes, and Nobre, 2017). Here, we use simultaneous multi-region population recordings to understand how this selection process transforms the contents of working memory in the service of behavior.

To investigate the neural basis of selection, we trained monkeys to perform a visual working memory task that required them to hold two items in working memory and then, after a cue, select the task-relevant item to guide a future behavioral decision. Using large-scale neural recordings in prefrontal, parietal, and visual cortex, we found that selection increases the amount of information about the task-relevant item in neural firing rates. Furthermore, selection transformed the geometry of population codes for the remembered items. Before selection, when both items were relevant to the task, the identity of each item was represented in an independent subspace of neural activity. After selection, item population codes were transformed, with the selected item moving into a new subspace that encoded information relevant for the decision at the end of the trial while ‘marginalizing out’ information about now-irrelevant stimulus dimensions. Together, our results provide novel insight into the neural mechanisms that control the contents of working memory.

1-025. Geometric learning enables 3D kinematic profiling across species

Tim Dunn[1] TIMOTHY.DUNN@DUKE.EDU
Kyle Severson[1] KYLE.SEVERSON@DUKE.EDU
Diego Aldarondo[2] DIEGOALDARONDO@G.HARVARD.EDU
David Hildebrand[3] DAVID@HILDEBRAND.NAME
Selmaan Chettih[4] SELMAANCHETTIH@GMAIL.COM
William Wang[2] WLWANG70@GMAIL.COM
Amanda Gellis[2] AJG1GELLIS@GMAIL.COM
David Carlson[1] DAVID.CARLSON@DUKE.EDU
Dmitriy Aronov[5] DA2006@COLUMBIA.EDU
Winrich Freiwald[3] WINRICH.FREIWALD@ROCKEFELLER.EDU
Fan Wang[1] FAN.WANG@DUKE.EDU
Bence Olveczky[2] EMBERBENCE@GMAIL.COM
Jesse Marshall[2] JESSE.D.MARSHALL@GMAIL.COM

[1]Duke University
[2]Harvard University
[3]The Rockefeller University
[4]Columbia University
[5]Zuckerman Mind Brain Behavior Institute, Columbia University

Understanding the neural computations that give rise to complex, ethological behaviors requires an ability to precisely measure freely moving animal behavior in naturalistic environments. While tracking with 2D convolutional neural networks (CNNs) has improved our ability to monitor behavioral correlates in confined tasks, these approaches scale poorly to 3D tracking in freely moving animals, where appearance changes and occlusions are ubiquitous. To address this, we developed DANNCE, a 3D CNN inspired by recent advances in geometric deep learning, that enables whole-body 3D kinematic recordings in naturalistic environments. DANNCE uses projective geometry to construct a 3D metric feature space in which multiple camera views can share information and 3D priors for animal pose can be learned. To train DANNCE, we collected Rat7M, a new 7-million frame video and motion capture dataset. Rat 7M will be the first 3D animal motion capture dataset available to the animal tracking community, and we hope that it will fuel algorithm development for model system neurobiology, similar to how 3D human motion capture datasets have transformed algorithm development in deep computer vision. Once trained, DANNCE takes videos from as few as 1 camera and outputs 3D coordinates of pre-specified body landmarks. We demonstrate that DANNCE tracks rats and mice with orders of magnitude greater sensitivity than DeepLabCut (DLC), and we used DANNCE to track whole-body 3D kinematics over hours-long recording sessions in rats, mice, marmosets, and chickadees. We show how DANNCE can map the behavioral repertoire of these organisms and kinematically profile behaviors, such as grooming, that have been difficult to assess with hand annotation. DANNCE offers unprecedented analytical access to behavior.

1-026. Cortico-cerebellar networks as brain-wide decoupling machines

Ellen Boven[1,2] ELLEN.BOVEN@BRISTOL.AC.UK
Joseph Pemberton[1] OQ19042@BRISTOL.AC.UK
Richard Apps[1] R.APPS@BRISTOL.AC.UK
Rui Ponte Costa[1] RUI.COSTA@BRISTOL.AC.UK

[1]University of Bristol
[2]Physiology Pharmacology and Neuroscience

The cerebellum has widespread connections with the cerebral cortex. However, it is not well understood what is the function of these cortico-cerebellar networks. Here we propose that the cerebellum provides the cerebral cortex with approximations of forward and feedback cortical computations. Such approximations can then be used to speed up forward and/or feedback computations by decoupling them from the rest of the network. Inspired by a recent deep learning algorithm, decoupled neural interfaces (DNI), we introduce a cortico-cerebellar modelling framework that approximates cortical computations – cortico-cerebellar DNI (CC-DNI). To illustrate the potential of this framework we focus on modelling a given brain area as a recurrent neural network in which the cerebellum approximates temporal feedback signals. First, we show that this CC-DNI mechanism facilitates learning in a range of sensorimotor tasks. Second, testing CC-DNI in a cognitive task, i.e. caption generation, demonstrates that our model readily applies to a wider range of modalities. Models without the cerebellar-DNI component exhibit deficits similar to what has been observed in cerebellar patients in both motor and cognitive tasks. Finally, we introduce a CC-DNI model with highly sparse connectivity as observed in the cerebellum, which improves the model efficiency and learning through decorrelation. Overall, our work offers a novel perspective on the cerebellum as a brain-wide decoupling machine for efficient cortical processing and a new avenue of research between deep learning and neuroscience.

1-027. Grid cell path integration for movement-based visual object recognition

Niels Leadholm NIELS.LEADHOLM@SEH.OX.AC.UK
Marcus Lewis MLEWIS@NUMENTA.COM
Subutai Ahmad SAHMAD@NUMENTA.COM

Numenta

Grid cells enable the brain to model the physical space of the world and navigate effectively via path integration, by updating global self-position using self-movement. Recent proposals suggest that the brain uses similar mechanisms to understand the structure of objects in diverse sensory modalities (Hawkins et al, 2019), including vision (Bicanski & Burgess, 2019). In machine vision, visual object recognition given a sequence of sensory samples of an image, such as saccades, is a challenging problem when the sequence does not follow a consistent, fixed pattern. Yet this is something humans do naturally and effortlessly. In this submission we explore how grid cell-based path integration in a cortical network can reliably recognize objects given an arbitrary sequence of inputs. Our network (GridCellNet) uses grid cell computations to integrate visual information and make predictions based on upcoming movements. We use local Hebbian plasticity rules to learn rapidly from a handful of examples (few-shot learning), and consider the task of recognizing MNIST digits given only a sequence of image feature patches. We compare GridCellNet to k-Nearest Neighbour classifiers as well as supervised recurrent neural networks (RNNs), both of which lack explicit mechanisms for handling arbitrary sequences of input samples. We show that GridCellNet can reliably perform classification, and generalize to both unseen examples and completely novel sequence trajectories. We further show that inference is often successful after sampling a fraction of the input space, enabling the predictive network to quickly reconstruct the rest of the image given just a few movements. We propose that dynamically moving agents with active sensors can use grid cell representations for efficient recognition and prediction, as well as navigation.

1-028. Recurrent networks fitting neural temporal responses to natural images exhibit contextual modulation

Harold Rockwell[1,2] HROCKWEL@ANDREW.CMU.EDU
Yimeng Zhang[1] YIMENGZH@ANDREW.CMU.EDU
Gaya Mohankumar[1] GAYATHRIMOHANKUMAR95@GMAIL.COM
Stephen Tsou[1] STSOU@ANDREW.CMU.EDU
Tai Sing Lee[1] TAISLEE@ANDREW.CMU.EDU

[1]Carnegie Mellon University
[2]Center for the Neural Basis of Cognition

We propose a convolutional neural network with a layer of lateral recurrent connections to predict the observed full temporal responses of a neuronal population. The bottom-up receptive fields are obtained through transfer learning from an intermediate layer of a standard Imagenet-trained DenseNet-121 and the recurrent kernels are learned to fit measured neuronal responses, producing high predictive performance. We performed two standard neurophysiological V1 experiments on the hidden units and found that the hidden units exhibit quintessential contextual modulation effects observed in V1, namely longitudinal facilitation and lateral suppression of oriented bars as in association field (Kapadia et. al., 1999), as well as contextual modulation resulting in dynamic reduction in orientation bandwidths and spatial frequency bandwidths over time (Ringach et. al., 1997, 2002). Our results demonstrate that deep learning models with appropriately structured recurrent circuits, trained end-to-end for neural response prediction, can be meaningfully analyzed and reproduce neurophysiological phenomena, therefore potentially providing a computational approach to investigate the mechanisms and circuits of early visual cortex.

1-029. Data-driven model of auditory thalamocortical system rhythms

Erica Griffith[1,2] ERICA.GRIFFITH@DOWNSTATE.EDU
Salvador Dura-Bernal[3] SALVADORDURA@GMAIL.COM
Annamaria Barczak[4] ANNIEBARCZAK@GMAIL.COM
Monica O'Connell[4] OCONNELLEQUAL@GMAIL.COM
Tammy McGinnis[5] TMCGINNIS@NKI.RFMH.ORG
William Lytton[6] BILL.LYTTON@DOWNSTATE.EDU
Peter Lakatos[7] PLAKATOS@LAKATOSLAB.COM
Samuel Neymotin[7] SAMNEMO@GMAIL.COM

[1]SUNY Downstate Medical Center
[2]Neural and Behavioral Sciences
[3]SUNY Downstate Medical Center, Nathan S. Kline Institute for Psychiatric Research
[4]Nathan S. Kline Institute for Psychiatric Research
[5]Nathan S. Kline Institute for Psychiatric Research, SUNY Downstate Medical Center
[6]SUNY Downstate Medical Center, Kings County Hospital Center
[7]Nathan S. Kline Institute for Psychiatric Research, NYU Langone Medical Center

Entrainment is the alignment of low-frequency cortical oscillation phase to the phase of incoming auditory stimuli. Speech waveforms in particular have an oscillatory, quasi-rhythmic quality. The cochlea filters these raw waveforms according to frequency, resulting in filtered signals whose envelopes retain this rhythmicity. The frequency of these envelopes lies reliably in the 4-7 Hz range, even across languages. This frequency range matches the rate at which syllables occur in continuous speech, and furthermore, it has been repeatedly observed that low-frequency oscillations in auditory cortex phase-lock to the envelopes of these filtered signals. These observations have led to a growing body of data showing that phase-locking between cortex and speech waveform is a crucial component to the parsing and comprehension of continuous speech, with disruptions in entrainment reliably leading to reduced speech intelligibility. However, despite the growing body of behavioral work investigating the functions and significance of entrainment, the cellular and network-level mechanisms underlying this phenomenon are not fully understood. This project thus uses a dual in vivo and in silico approach to investigate the mechanisms underlying the oscillatory cortical activity seen during speech processing. The in vivo arm of this study gathered invasive electrophysiology recordings from superior temporal gyrus in human subjects, and primary auditory cortex in nonhuman primates, illustrating what oscillatory activity was present at rest, and how these oscillations changed in response to stimuli. The in silico branch of this study used this experimental data to construct and tune a biophysically detailed model of the macaque auditory thalamocortical system. The model's detailed parameters are being used to investigate specific thalamocortical circuit hypotheses regarding the mechanisms of entrainment: for example, the idea that entrainment relies on a "phase reset" mechanism where supragranular interneurons, activated by thalamic matrix inputs, contribute inhibition which sculpts the signal in lower cortical layers.

1-030. Intracortical inhibition is precisely organized in early cortical networks

Gordon Smith[1,2] GBSMITH@UMN.EDU
Bettina Hein[3] BH2757@COLUMBIA.EDU
Matthias Kaschube[4] KASCHUBE@FIAS.UNI-FRANKFURT.DE
Haleigh Mulholland[1] MULHO042@UMN.EDU

[1]University of Minnesota
[2]Neuroscience
[3]Columbia University
[4]Frankfurt Institute for Advanced Studies

A fundamental feature of the visual cortex in primates and carnivores is the presence of columnar functional networks which span millimeters of cortical area. Recent work in the ferret demonstrates that these functional networks are present in correlated spontaneous activity over 10 days before eye opening and well before the maturation of long-range horizontal connectivity. Neural field models indicate that such long-range correlations can be generated entirely by local intracortical circuitry, however such models typically display spatially co-aligned patterns of excitatory and inhibitory activity. To date, it remains unclear whether intracortical inhibition in the early cortex exhibits the fine scale modular organization required for these models or rather displays a broad and diffuse spatial structure, pooling activity over large spatial areas. Here, we use in vivo calcium imaging in specifically labeled excitatory and inhibitory neurons to directly investigate the spatial organization of intracortical inhibition and its relationship to excitatory activity in the early visual cortex. We first demonstrate that the developmental switch from excitatory to inhibitory GABAergic signaling has occurred by the age in which large-scale functional networks are apparent, fulfilling a critical requirement of published models. We find that inhibitory neurons exhibit precise functional organization in the early cortex, showing highly modular activity and long-range correlations with a strength and spatial scale similar to that found in excitatory cells. Strikingly, inhibitory networks display robust fine-scale organization, showing abrupt discontinuities ('fractures') in network organization across the cortex. Finally, using simultaneous 2-photon imaging, we show that excitatory and inhibitory activity is balanced on a local spatial scale and exhibit precisely matched patterns of correlated spontaneous activity. These findings clearly demonstrate a fine-scale organization of intracortical inhibition already in the developing cortex, and support the ability of large-scale cortical networks to self-organize through precisely correlated local excitatory and inhibitory activity.

1-031. Inference of mouse V1 connectivity from visual responses using stabilized supralinear network model

Nataliya Kraynyukova[1] NATALIYA.KRAYNYUKOVA@BRAIN.MPG.DE
Simon Renner[2] RENNER.SIMON@CAMPUS.LMU.DE
Gregory Born[2] BORN@BIO.LMU.DE
Yannik Bauer[2] Y.BAUER@BIOLOGIE.UNI-MUENCHEN.DE
Ann H. Kotkat[3] ANN.KOTKAT@TUM.DE
Xinyu Liu[2] XINYU.LIU@CAMPUS.LMU.DE
Martin Spacek[2] M.SPACEK@LMU.DE
Georgi Tushev[4] GEORGI.TUSHEV@BRAIN.MPG.DE
Laura Busse[2] BUSSE@BIO.LMU.DE
Tatjana Tchumatchenko[4] TATJANA.TCHUMATCHENKO@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]LMU Munich
[3]TU Munich
[4]MPI for Brain Research

Recent experimental work provides detailed functional and structural connectivity diagrams of cortical circuits. However, the relation between experimentally reconstructed connectomes and the connectivity constraints derived from recorded network responses remains largely unclear. Here, we infer the synaptic wiring pattern of mouse primary visual cortex (V1) from in-vivo recordings of contrast and orientation responses in V1 and dLGN using a stabilized supralinear network model (SSN). The SSN is a recurrent network model consistent with such fundamental cortical computations as normalization, surround suppression, and contrast-invariance. Our inference method is based on the property of contrast-invariance, meaning that the width of orientation tuning curves does not considerably change with contrast. We confirm contrast-invariance experimentally in V1 neurons and show that it can serve as a computational constraint shaping synaptic wiring. The inferred V1 connectivity and dLGN input weights share many features with previously measured synaptic connectivities. For example, we find that the excitatory-to-excitatory connection is weaker than the inhibitory-to-inhibitory, and the inhibitory external input weight is stronger than the external excitatory input. In conclusion, our SSN-based inference method links

recorded activity in the intact brain with a range of connectivity and input distributions consistent with neural wiring statistics provided by experimental connectomics studies.

1-032. Manifold attractors without symmetries: Spatial correlations shape temporal dynamics in neural manifolds

Alexander Rivkind[1,2] SASHKARIVKIND@GMAIL.COM
Ran Darshan[3] DARSHANR@JANELIA.HHMI.ORG

[1]Weizmann Institute of Science
[2]Department of Neurobiology
[3]Howard Hughes Medical Institute

Computations by manifold attractors are the standard framework for understanding neural representations and computations involving continuous variables. This framework is the dominant paradigm, and often the only one, in modeling various systems in neuroscience, e.g. parametric and spatial working memory or path integration. The main assumption in this theory is a symmetry property in the synaptic connections. For example, in the classical ring attractor model for representing and updating of head direction, a rotation symmetry is assumed: connectivity in the network depends solely on the heading preferences of the neurons. This assumption has been debated for years. In the presence of slightest alternations in the network connectivity, the symmetry breaks and the computational capabilities of such networks quickly deteriorate as a result of a drift towards a few isolated attractors. It is thus unclear if the concept of computations by manifold attractors applies to real biological systems in which imperfections are inevitable. Here we show that such symmetries are not necessary and manifold attractors can emerge in recurrent networks even without such symmetries. We present a minimal and solvable model of a trained recurrent neural network which stores and updates the representation of a continuous variable, but lacks any symmetry. As a result, we show how manifold attractors can cope with diverse neuronal responses and a high level of synaptic heterogeneities. Based on mean-field analysis of the model we find that a manifold attractor emerges remarkably quickly with the number of training samples. We further analyze the computational capabilities of the trained network, showing that the dynamics in the vicinity of the attractor is controlled by a few leading principal components of neural activity. Our framework connects the spatial correlations of neural activity to the build-up of manifold attractors and to its dynamical properties. It suggests a general principle for how spatial correlations predict the dynamics in putative manifold attractors in the brain.

1-033. Predicting real time behavior changes using unsupervised learning models

Chengxu Zhuang[1] CHENGXUZ@STANFORD.EDU
Xiaoxuan Jia[2] XIAOXUANJ@ALLENINSTITUTE.ORG
Violet Xiang[1] ZIYXIANG@STANFORD.EDU
Yoon Bai[3] YHB@MIT.EDU
James DiCarlo[3] DICARLO@MIT.EDU
Daniel Yamins[1] YAMINS@STANFORD.EDU

[1]Stanford University
[2]Allen Institute
[3]Massachusetts Institute of Technology

Deep convolutional neural networks (DCNNs) have provided the best quantitative models of the neural responses throughout the ventral visual stream (VVS). Recent works have even shown that such models can be yielded from DCNNs trained using contrastive learning methods, a family of newly developed and high-performant unsupervised learning algorithms, making these DCNNs potential candidates for biologically-plausible computational models for the VVS. However, these models are so far only compared to the neural responses averaged across trials in multiple sessions, despite the fact that even the neural representations in the adult VVS dynamically reorganize as an online-learning effect from the presented stimulus and the resulting changes yield meaningful and observable behavior changes. Modeling these online-learning effects of the VVS is therefore critical to fully modeling the VVS. In this work, we leverage the best unsupervised DCNN models of the VVS and the corresponding unsupervised learning algorithms to build computational models to predict the real time categorization behavior changes measured from human subjects. We first pretrain the DCNN on naturalistic images using unsupervised learning algorithms and then present the network the same data stream of images that a series of human subjects have seen. Finally, we compare the changes in categorization performance on specific sets of images between the humans and the models. We find that the qualitative changes in the human and the model are the same. More specifically, presenting pairs of images belonging to the same category but in different size will strengthen size-

tolerance towards this category, while presenting pairs of images belonging to different categories will weaken size-tolerance in both categories. However, a quantitative mismatch still exists between the learning effects in models and humans, suggesting that more powerful learning algorithms are needed.

1-034. Sensitivity to higher-order texture statistics in segmentation reflects redundancy in natural images

Daniel Herrera-Esposito[1,2] DHERRERA1911@GMAIL.COM
Leonel Gomez-Sena[1] LEONEL.GOMEZ@GMAIL.COM
Ruben Coen-Cagli[3] RUBEN.COEN-CAGLI@EINSTEINMED.ORG

[1]Universidad de la Republica
[2]Biology Department
[3]Albert Einstein College Of Medicine

Segmenting the visual input is essential for parsing visual scenes into objects and surfaces. This process relies on segmentation cues such as disparity, color, motion and texture. Although texture cues have been extensively studied, most studies have focused on the role of second-order texture statistics (spectral statistics), despite the known perceptual relevance of higher-order statistics (HOS). In particular, HOS of the influential Portilla-Simoncelli texture model have been widely used to describe neural selectivity in visual areas V2 and V4, and to unify several disparate perceptual phenomena, but their role as a segmentation cue is uncertain. In this work we used a psychophysical task to test whether differences in the Portilla-Simoncelli HOS between two image regions support segmentation in humans, and to compare their effect to spectral statistics. We show that these HOS induced weak segmentation in human subjects, while spectral statistics induced strong segmentation. This is surprising considering the known relevance of these HOS for perception and physiology. Thus, we next hypothesized that this unexpected result may reflect a property of natural images, namely, a redundancy between spectral and HOS for the task of segmentation. To test this, we trained linear and non-linear classifiers to solve a segmentation task on natural images, using either spectral statistics, HOS, or both together. We observed that while both sets of statistics alone allow for good segmentation performance, combining them produced weak improvements to performance, indicating a redundancy for natural image segmentation. Given resource constraints, this may lead the brain to rely mostly on spectral statistics. Although natural statistics are commonly used to explain perceptual phenomena, task-specific analysis of natural statistics are less common. Here we show how considering natural image statistics in a task-specific manner may explain variations in the relevance of different image features across perceptual tasks in biological vision.

1-035. Distinct higher-order representations of natural sounds in human and ferret auditory cortex

Agnes Landemard[1,2] AGNES.LANDEMARD@ENS.FR
Celian Bimbard[3] C.BIMBARD@UCL.AC.UK
Charlie Demene[4] CHARLIE.DEMENE@GMAIL.COM
Shihab Shamma[1,5] SAS@UMD.EDU
Sam Norman-Haignere[6] SN2776@COLUMBIA.EDU
Yves Boubenec[1,7] YVES.BOUBENEC@ENS.FR

[1]Ecole Normale Superieure
[2]Departement d'etudes Cognitives
[3]University College London
[4]ESPCI Paris
[5]University of Maryland
[6]Zuckerman Mind Brain Behavior Institute, Columbia University
[7]Cognitive Studies

Little is known about how neural representations of natural sounds differ across species. For example, speech and music play a unique role in human hearing, yet it is unclear how auditory representations of speech and music differ between humans and other animals. Using functional Ultrasound imaging, we measured responses in ferrets to a set of natural and spectrotemporally-matched synthetic sounds previously tested in humans. Ferrets showed similar lower-level frequency and modulation tuning to that observed in humans. But while humans showed prominent selectivity for natural vs. synthetic speech and music in non-primary regions, ferret responses to natural and synthetic sounds were closely matched throughout primary and non-primary auditory cortex, even when tested with ferret vocalizations. This finding reveals that auditory representations in humans and ferrets diverge sharply at late stages of cortical processing, potentially driven by higher-order processing demands in

speech and music.

1-036. Erasing motion: The scrambling of direction selectivity in visual cortex during saccades

Satoru Miura[1,2] SATMIURA14@GMAIL.COM
Massimo Scanziani[1] MASSIMO@UCSF.EDU

[1]University of California, San Francisco
[2]Department of Physiology

Sensory stimuli are often generated by the animal's own movements. Many sensory systems suppress these self-induced inputs and selectively respond to externally generated stimuli, thereby maintaining the perceptual stability of the external world. Rapid eye movement, called saccades, induces motion of the visual scene on the retina. Several mechanisms have been proposed that prevent the perception of such motion. Yet, how they are implemented in the circuitry of visual cortex is poorly understood. Here, we focused on the mouse primary visual cortex (V1). In V1, neurons encode the motion direction of visual stimuli. Using extracellular recording, we began by comparing the response of V1 neurons to spontaneous saccades and their response to external visual stimuli designed to mimic the motion of visual scene induced by saccades, which we termed "pseudo-saccades". Many neurons showed direction selectivity for saccades (i.e., responded more to saccades moving in one direction than to those in the other direction) as well as for pseudo-saccades. Interestingly, however, there was no correlation between the direction selectivity for the two conditions. Consequently, a classifier that discriminates the direction of externally induced visual motion based on the population activity in V1 is unable to discriminate the direction of motion induced by actual saccades. Thus, in V1, direction selectivity is scrambled during saccades, and information about the direction of motion is suppressed. Where does this scrambling originate from? We discovered that V1 receives a non-visual input from the pulvinar nucleus of the thalamus around the time of saccades. To determine whether this pulvinar input mediates the scrambling of direction selectivity, we pharmacologically silenced the pulvinar. Upon silencing this thalamic nucleus, the direction selectivity for real saccades and pseudo-saccades became correlated. Thus, the pulvinar erases saccade-induced motion of the visual scene in V1 by scrambling direction selectivity.

1-037. The role of compositional abstraction in human and artificial neural networks

Takuya Ito[1,2] TAKU.ITO1@GMAIL.COM
Tim Klinger[3] TKLINGER@US.IBM.COM
Douglas Schultz[4] DHSCHULTZ29@GMAIL.COM
Michael Cole[1] MICHAEL.COLE@RUTGERS.EDU
Mattia Rigotti[3] MR2666@COLUMBIA.EDU

[1]Rutgers University
[2]Center for Molecular and Behavioral Neuroscience
[3]IBM Research AI
[4]University of Nebraska-Lincoln

Humans have a remarkable ability to rapidly generalize and learn new tasks. This characteristic is exemplified by compositional reasoning – the ability to rapidly reuse previously learned concepts and rules in new contexts. In this work we illustrate how compositional abstraction facilitates generalization to new task contexts in whole-brain human fMRI data and artificial neural networks (ANNs) during a 64-context sensorimotor task. First, we identified behavioral correlates of compositional reasoning during task performance in humans. This was illustrated by higher performance on unseen tasks containing more overlapping rules with previously learned tasks. Second, using recently characterized measures of representational abstraction, we demonstrated the existence of abstract representational geometries tiled across human cortex. Interestingly, these abstract representations differed in space and magnitude across distinct compositional domains (such as logical, sensory, and motor gating), suggesting that the brain abstracts distinct types of information differently. Third, in ANNs, we showed that compositional training on primitive tasks and concepts facilitates generalization and zero-shot performance on novel compositional tasks. Specifically, training on primitive tasks enabled systematicity in ANNs – the ability to systematically re-use previously learned concepts on unseen task sets. Furthermore, compositional learning promoted abstract hidden representations, suggesting that successful generalization to new contexts requires abstract representations. Together, our findings provide empirical and computational evidence of the role of abstraction in compositional generalization in human and artificial neural networks.

1-038. Information-limiting correlations enhance information transmission in biophysical neuronal models

Giulio Bondanelli[1,2] GIULIO.BONDANELLI@IIT.IT
Monica Moroni[1] MONICA.MORONI@IIT.IT
Christopher Harvey[3] CHRISTOPHER_HARVEY@HMS.HARVARD.EDU
Stefano Panzeri[1] STEFANO.PANZERI@IIT.IT

[1]Istituto Italiano di Tecnologia
[2]Neural Computation Laboratory
[3]Harvard Medical School

Upon multiple stimulus presentations, neurons exhibit trial-to-trial correlated variability. A number of experimental and theoretical works have established that, at the population level, correlations aligned with the stimulus dimensions limit stimulus information [1], leading to the widely-held surmise that information-limiting correlations are detrimental for behavioural discrimination. However, contrarily to this expectation, recent studies have reported that correlations are stronger in trials where the animal correctly performed perceptual discrimination tasks than in error trials [2,3,4]. Recent experimental evidence collected from mouse posterior parietal cortex (PPC) during perceptual tasks has suggested that these results can be reconciled if the downstream behavioural readout of neural population activity is enhanced by correlations [3]. To date, how downstream neurons may use correlations to overcome their information-limiting effect remains an open theoretical question. Here we test the hypothesis that correlations may be integrated at the single-neuron level to overcome their information-limiting effect. We propose a minimal biophysical model consisting of a readout neuron integrating presynaptic correlated inputs. We measured the amount of stimulus information transmitted from the input to the readout neuron as a function of biophysical parameters and input correlations. We found that, in the regime where the average input to the readout is below spiking threshold, correlations substantially enhance the transmission of information through the network even though they hinder input information. Lastly, we directly compared the model to PPC data from [3]. As in PPC data, the model predicted higher correlations in correct than in error trials. Further, fitting different logistic readout models to the output activity revealed that correlations were used to amplify the transmission of stimulus information to the readout neuron. Together, these findings suggest that a single-neuron mechanism may be sufficient for the behavioural readout of association cortex to counteract the negative effect of correlations on the encoded stimulus information.

1-039. Chaotic dynamics in a spatially distributed network generate population-wide shared variability

Noga Mosheiff[1,2] NOW17@PITT.EDU
Bard Ermentrout[1] BARD@PITT.EDU
Chengcheng Huang[1,2] HUANGC@PITT.EDU

[1]University of Pittsburgh
[2]Neuroscience

Neural activity in the cortex is highly variable in response to an identical stimulus. This variability is correlated among neurons. The correlation has been shown to decrease with the distance between a pair of neurons. Recently, cortical population recordings revealed that the shared variability in a neural population is also low dimensional. However, the source of the population-wide shared variability is unknown. Previous work (Huang et al. 2019) demonstrated that spatiotemporal dynamics in spiking neural networks can give rise to population-wide shared variability. In this work, we show that population-wide shared variability can be internally generated in deterministic recurrent networks with chaotic rate dynamics. We analyzed the dynamical regimes of recurrent networks with excitatory and inhibitory populations, which are spatially distributed on a ring or a torus. We found that the networks have dynamical regimes of different regular solutions, depending on the temporal and spatial scales of inhibition, such as static bumps, bulk oscillations, traveling waves and alternating bumps. In particular, in a parameter regime where the excitatory projection width is close to the inhibitory projection width, the networks exhibit chaotic rate dynamics. The chaotic solutions contain low-frequency power in rate variability and have low-dimensional and distance-dependent correlations, in agreement with experimental findings. Further, we found that chaos can be induced by correlated noisy inputs. The chaotic parameter regime expands as the variance of the correlated noise component increases, while being insensitive to the variance of the independent noise components.

1-040. The excitability-functionality trade-off: Random graph models of epilepsy

Zhenrui Liao[1,2] ZHENRUI.LIAO@COLUMBIA.EDU
Attila Losonczy[1] AL2856@COLUMBIA.EDU
Christos Papadimitriou[1] CHRISTOS@COLUMBIA.EDU

[1]Columbia University
[2]Department of Neuroscience

Epilepsy is a network pathology in which circuits involved in normal brain function are recruited to deleteriously high oscillatory network activity. Here we investigate whether epilepsy, seen this way, is a con- sequence of a more general network phenomenon rooted in the inherent randomness of synaptic connections. We study two simple models of local network connectivity based on the classical Erdős–Renyi model of random graphs, and establish in both a trade-off between functionality and the risk of unchecked activity. We show in each of the two models that, if a key parameter of the model is large enough to guarantee functionality of the network, then the network is also very likely to have the capacity for unbounded oscillation.

1-041. Dopamine's function in the algorithmic basis of foraging decisions

Sarah Starosta[1,2] SARAH.STAROSTA@WUSTL.EDU
Sergey Shuvaev[3] SSHUVAEV@CSHL.EDU
Duda Duda Kvitsiani[4] KVITSI@DANDRITE.AU.DK
Alexei Koulakov[3] KOULAKOV@CSHL.EDU
Adam Kepecs[5] AKEPECS@WUSTL.EDU

[1]Washington University In St. Louis
[2]Neuroscience
[3]Cold Spring Harbor Laboratory
[4]Aarhus University
[5]Washington University in St. Louis

We are continually confronted with decisions about whether to stay engaged with the current option or to switch to a new one. These stay-or-leave decisions include finding a job, buying a house, and selecting a partner. They have been usually studied in ecology as foraging decisions. However, little is known about their neural basis. Thus, we set out to study the behavioral algorithm and neural mechanisms underlying stay-or-leave decisions in a behavioral paradigm inspired by foraging theory. In this paradigm, mice were facing the decision when to leave depleting reward sources. To explore the trial-by-trial choice strategy and its neural correlates, we implemented several reward manipulations and performed optical recordings of dopamine neuron activity in the Ventral Tegmental Area (VTA) – a brain area involved in reward-guided behaviors (Schultz et al, 2015). We compared the observed mouse behavior to different computational models and identified a novel decision rule that accounts for animals' behavior in a range of conditions, including different inter-trial intervals, reward depletion rates, and average reward rate of the environment. Specifically, we found that the animals compare the next expected reward to the exponential average of the previous rewards – a decision rule we named the leaky MVT, for some similarity with the conclusions of the Marginal Value Theorem (MVT; Charnov, 1976; Constantino & Daw, 2015). We show this decision rule may be learned via a reinforcement learning (RL) paradigm called R-learning (Schwartz, 1993), but is not consistent with classical V-, or Q-learning paradigms as these make qualitatively different predictions. Finally, we show that dopaminergic signaling in the VTA best correlates with the Reward Prediction Error of R-learning, pointing to a potential learning mechanism that optimizes stay-or-leave choices. Overall, our work offers an algorithmic decision rule and neuronal implementation for an ethologically relevant behavior based on qualitative different model predictions.

1-042. The Jacobian switching linear dynamical system

James Smith[1,2] JSMITH14@STANFORD.EDU
Scott Linderman[1] SCOTT.LINDERMAN@STANFORD.EDU
David Sussillo[3] SUSSILLO@GOOGLE.COM

[1]Stanford University
[2]Institute for Computational and Mathematical Engineering
[3]Google / Stanford

We are interested in models that improve our understanding of how nonlinear, recurrent neural networks (RNNs) compute since they are broadly used in neuroscience [1,2]. Recent work has shown that switching linear dy-

namical systems (SLDS) with discrete and continuous states can reasonably approximate nonlinear dynamical systems [3,4]. SLDSs are desirable models since they potentially provide flexible modeling capacity while retaining some interpretability. However, SLDSs come with challenges such as the need to manually choose the number of discrete states and having to solve the difficult combinatorial optimization problem of assigning those states to the various portions of the dynamics. Further, nonlinear RNNs often have continuous manifolds of fixed points, which suggests a finite number of discrete linear regimes may be inadequate.

Here we propose a novel model called the Jacobian SLDS (JSLDS), taking inspiration from fixed point analysis of nonlinear RNNs [5]. In particular, the JSLDS is an SLDS that is based on the Jacobian of a separate nonlinear RNN it is co-trained with. JSLDS learns to approximate a nonlinear RNN by co-training with the nonlinear RNN while learning its fixed points and when to switch around them. By training a JSLDS, we obtain a first-order model of the nonlinear RNN we have just trained. This generalizes the standard SLDS to a continuous mixture of linear dynamical systems, each specified by the Jacobian of an RNN about a point in state space. The JSLDS confers advantages over the standard SLDS: it sidesteps the discrete optimization problem by allowing a continuum of switching states, namely the fixed points of the RNN with which it is co-trained. Finally, the JSLDS can be used in multiple settings: as a more interpretable task-based model than generic RNNs, and as a flexible data-model for neural and/or behavioral recordings.

1-043. Seizure detection and localization using human intracranial electrophysiology via information theoretic methods

Lisa Yamada LYAMADA@STANFORD.EDU
Tomiko Oskotsky TNGUYEN9@BILAB.US
Paul Nuyujukian 21COSYNE@PN.STANFORD.EDU

Stanford University

Of the 1% of the world population with epilepsy, one-third have refractory epilepsy, in which their only option to manage seizures is a high-risk surgery to remove seizure onset zones (SOZs, brain regions that are most likely to cause seizures). The current state of epilepsy treatment heavily relies on manual evaluation of EEG by epileptologists and lacks interventions that leverage the rich information in EEG. To combat this limitation, we introduce an information theoretic estimate of joint entropy called inverse compression ratio (ICR), which utilizes compression algorithms, as a measure of quantitative EEG (qEEG). With our data repository of continuous, 10kHz intracranial EEG acquired from clinical neuromonitoring studies of adult and pediatric participants, we study the relationship between ICR and seizure activity. When comparing ICR across time, we observed a sharp peak in ICR at seizure onset, followed by a dip before returning to baseline. Furthermore, when analyzing characteristics of the ICR peak at seizure onset (e.g. peak amplitude) across neural channels, we observed prominent changes that distinguish channels located in SOZs. When using ICR to perform seizure detection, we found a sensitivity/specificity (SE/SP) of 86%/98.3% across 5 participants (30 seizures). In comparison, variance had a SE/SP of 44.7%/96.4%. Our results demonstrate that information theoretic measures of ICR performs comparably, if not better, than other qEEG measures, suggesting their potential in seizure detection and localization. Previous studies on entropy of EEG may not have detected the brief spike in information content at seizure onset due to signal quality or limitations in sampling rate. Implementing robust qEEG techniques to clinical practice may offload labor-intensive tasks from clinicians and uncover EEG features that cannot be detected by eye, broadening our understanding of epilepsy and improving therapy.

1-044. Synaptic plasticity shapes top-down modulation of cortical circuits with multiple interneuron types

Marcel Juengling[1] MARCEL.JUENGLING@BRAIN.MPG.DE
Leonidas Richter[2] LEONIDAS.RICHTER@BRAIN.MPG.DE
Luca Mazzucato[3] LMAZZUCA@UOREGON.EDU
Julijana Gjorgjieva[1,4] GJORGJIEVA@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]Max Planck Institute for Brain Research; Technical University of Munich
[3]Institute of Neuroscience, University of Oregon, USA
[4]Computation in Neural Circuits Group

Flexible behavioral responses are crucial for adapting to a changing environment and an animal's own internal state. On the neural circuit level, this is implemented via modulations of neural activity in response to top-down inputs. This occurs as early as in the sensory cortex where projections to VIP-expressing interneurons play a vital role in the state-dependent modulation of neural activity. VIP neurons are part of a cortical microcircuit comprising pyramidal excitatory (E) and other inhibitory cell types such as PV- and SST-expressing interneurons. Previous work has demonstrated the intricate neural dynamics these circuits produce in response to top-down inputs. However, the neural mechanisms responsible for generating and controlling the spatiotemporal dynamics of these responses remain elusive.

Here, we elucidate how top-down modulations of neural responses naturally emerge in cortical circuits via an interplay of synaptic plasticity, bottom-up sensory input, and top-down inputs. We present a novel theoretical framework based on recurrent spiking networks with multiple inhibitory cell types and synaptic plasticity. Combining network simulations and mean-field theory, we show that the formation of neuronal assemblies shaped by Spike-timing-dependent plasticity and heterosynaptic normalization leads to attractor dynamics with spontaneous assembly reactivations. We further demonstrate that top-down inputs targeting VIP neurons alter the temporal dynamics of these reactivations which may elucidate state-dependent modulations of sensory information processing speed. When inhibitory plasticity only drives potentiation of synapses from PV to E neurons, top-down modulation lengthens the timescale of reactivations. When synapses of SST to E neurons are also potentiated, top-down modulation shortens the timescale of reactivations.

Our work provides the first comprehensive study of the emergence and modulation of attractor dynamics in cortical circuits with multiple inhibitory cell types and emphasizes the role of synaptic plasticity in shaping network dynamics. Furthermore, our theory provides a testbed for cell-type-specific manipulation experiments.

1-045. A feedforward model accounts for the effects of masking in mouse V1

Dylan Barbera[1,2] DYLAN.BARBERA@UTEXAS.EDU
Nicholas Priebe[1] NICO@AUSTIN.UTEXAS.EDU
Lindsey Glickfeld[3] GLICKFELD@NEURO.DUKE.EDU

[1]University of Texas - Austin
[2]Institute for Neuroscience; Center for Learning and Memory
[3]Duke University

Neurons in the mouse primary visual cortex (V1) exhibit diverse responses to compound (plaid) stimuli. Unlike neurons in cat/primate V1, which are suppressed by a superimposed mask, many mouse V1 neurons have larger responses to plaids than to gratings (Juavinett & Callaway, 2015; Palagina et al., 2017; Rasmussen & Yonehara, 2017; Muir et al, 2015; Ringach et al., 2020). It is unknown whether a unique or shared mechanism is responsible for this diversity. Plaid stimuli have additionally been used to study motion integration in mouse V1. While some studies have demonstrated that responses in mouse V1 are selective for pattern (global) motion (Palagina et al., 2017; Rasmussen & Yonehara, 2017; Muir et al., 2015) others have reported the absence of this selectivity, similar to cat and primate V1 (Juavinett & Callaway, 2015). We aim to unify these lines of research by uncovering the mechanisms governing responses to plaids in mouse V1 using in-vivo electrophysiology and calcium imaging. In line with previous reports, we find interactions ranging from suppression to facilitation. Strongly orientation selective neuron responses exhibit cross-orientation suppression. Weakly selective neurons exhibit either facilitation or suppression. We demonstrate that we can systematically change whether a neuron is facilitated or suppressed by shifting the spatial phase of one of the plaid components. In addition, while there are some neurons that appear to exhibit selectivity for pattern motion, that selectivity is disrupted by changes in stimulus phase. These spatial phase dependent responses are a natural outcome of the spatial structure of the afferent receptive fields. Indeed, inactivation of cortex does not alter these response profiles and the data can be accounted for by a purely feedforward model. In sum, the diverse response of V1 neurons to plaid stimuli are easily accounted for by spatial structure of the thalamic afferents.

1-046. Texture-like representation for objects in ventral temporal cortex

Akshay Vivek Jagadeesh[1,2] AKSHAYJ@STANFORD.EDU
Justin Gardner[1] JLG@STANFORD.EDU

[1]Stanford University
[2]Department of Psychology

Does primate ventral temporal cortex (VTC) represent objects or textures? While VTC is widely thought to subserve object vision, it may do so by representing complex features that make up textures like gravel, bark, or fur, without specifically encoding the spatial configuration of features necessary to define an object. Here, we show that perceptually, humans are highly selective for object configuration, but deep convolutional neural networks (dCNNs), human VTC, and macaque inferotemporal (IT) cortex do not prioritize intact spatial configuration. That is, visual cortex has a texture-like representation, not a configurally-defined object representation. Unlike earlier methods (Lerner et al., 2001; Rust and DiCarlo, 2010) to probe object representation by scrambling stimuli which

disrupted both configuration and complex features, we developed a deep image synthesis algorithm (Freeman and Simoncelli, 2011; Gatys et al., 2015) to allow independent control over the complexity and configuration of features in synthesized images. Using an oddity detection task, we found that subjects' ability to detect intact objects degrades as synthesized decoy images have more complex matched features with more constrained configuration. Unlike humans, dCNN models' performance was near chance as feature complexity increased and paradoxically improved as feature configuration was more constrained. This suggests dCNNs are not selective for configuration, which we confirmed by showing that representational distance (correlation) between intact and scrambled objects was similar to that between two scrambled objects. Using BOLD imaging and a dataset of macaque IT neurons, we found that neural representations also lack selectivity for intact configuration. Our results suggest a misalignment between perception, which is highly selective for object configuration, and dCNN and VTC representations, which are configuration-invariant. We conclude that both dCNNs and VTC contain a texture-like representation of objects, and that further stages of neural computation are required to explain the configural selectivity of human object perception.

1-047. Coding of pattern complexity in V1 and V2 neurons

Corentin Massot[1,2] CORENTINMASSOT@YAHOO.FR
Xiaoqi Zhang[3] XIAOQIZH@ANDREW.CMU.EDU
Harold Rockwell[3,4] HROCKWEL@ANDREW.CMU.EDU
George Papandreou[5] GPAPAN@TTIC.EDU
Alan Yuille[6] AYUILLE@JHU.EDU
Tai Sing Lee[3] TAISLEE@CMU.EDU

[1]University of Pittsburgh
[2]Neurobiology
[3]Carnegie Mellon University
[4]Center for the Neural Basis of Cognition
[5]Toyota Technological Institute at Chicago
[6]Johns Hopkins University

V1 neurons are thought to encode simple features present in the visual information within their receptive fields (e.g. orientation and spatial frequency). However other features may have been overlooked due to the use of stimuli with simple shapes in past studies (Hegde, 2007; Tang, 2018). To overcome this issue, we designed a set of 50 stimuli that displayed varying shape complexity and average orientation. We recorded neurons in V1 and V2 of the macaque monkey using multi-contact laminar probes across cortical layers in single hypercolumns. First, we examined neuronal tuning across cortical regions. We found that the average response of all neurons displayed a clear preference for more complex stimuli. Tuning of individual neurons in V1 and V2 also showed an elevated average correlation (mean R2=0.37 (std 0.15) in V1 and mean R2=0.42 (std 0.18) in V2) with a complexity index. Second, we analyzed population responses across multiple sessions. We projected all stimuli into a low-dimensional space using the two first principal components of the response of all neurons in V1 and V2 or according to their laminar position. We found that the two first principal components were already sufficient to linearly encode both orientation and complexity. Regression coefficients showed that V1 population indeed explicitly encoded complexity (R2=0.47) though it encoded orientation better (R2=0.8). On the other hand, V2 population encoded complexity (R2=0.72) better than orientation (R2=0.4). Results were roughly similar across layers. Taken together, these results show that many neurons in V1 and V2 are tuned to the pattern complexity of the stimuli besides their orientation tuning at the individual level. The richer set of complex shapes used in this study revealed that there is a complexity axis at the population code level.

1-048. Beyond invariant object classification: new normative principles for models of IT visual cortex

Jack Lindsey[1,2] JACKWLINDSEY@GMAIL.COM
Elias Issa[1] ELIAS.ISSA@COLUMBIA.EDU

[1]Columbia University
[2]Neuroscience

Deep neural networks (DNNs) optimized for real-world object classification are the leading models predicting neural responses in inferior temporal cortex (IT). However, further optimizing DNN classification accuracy produces a saturating trend for predicting IT neural responses – potentially because pure performance optimization favors representations explicitly encoding information about object class at the expense of representing other sources of image-by-image variance. Here, we performed an extensive meta-analysis of current DNNs to identify represen-

tational properties underlying neural predictivity.

By examining an array of representational metrics – including classification performance, sparsity, and dimensionality – we identified two properties of DNN representations that were highly predictive of their matches to IT neural data: factorization of scene-to-scene variance from (1) viewpoint changes, induced by taking crops of an image or varying camera position in a video and (2) appearance transforms, induced by varying lighting and color. Factorizing (as opposed to being invariant to) scene viewpoint and appearance both matched or exceeded ImageNet classification accuracy in predicting the best models of high-level visual cortex across four datasets tested – two neural datasets in monkeys and two human fMRI datasets. Importantly, metric predictivity generalized across a diverse range of DNNs with varied architectures and objectives (n=47 models).

Consistent with these insights, we found that the models best matching neural data were self-supervised models optimized via contrastive learning for criteria similar to scene viewpoint and appearance factorization. These models improved upon the neural fits of architecture-matched controls trained for object classification. Based on our observations, we were able to simplify contrastive objective functions to bring them closer to biological plausibility while still yielding representations that predicted neural data well. Thus, our results revise the idea that IT is best explained through the lens of invariant object classification by suggesting new candidate normative principles guiding representations in high-level visual cortex.

1-049. Functional E-I spiking network with predictive coding

Veronika Koren[1,2] VERONIKA.KOREN@TU-BERLIN.DE
Tilo Schwalger[3,4] TILO.SCHWALGER@BCCN-BERLIN.DE

[1]Technische Universiatet Berlin
[2]Institutute of Mathematics
[3]TU Berlin
[4]Institute for Mathematics

One of the primary goals of neuroscience is to understand the computational principles that describe the formation of behaviorally relevant signals in the brain, as well as how are these computations realized within the constraints and properties of biological networks. Currently, most functional models of neural activity are based on firing rates, while the most relevant signals for inter-neuron communication are spikes. Recently, the framework of predictive coding has suggested a theory on how neural networks might compute behaviorally relevant signals with spikes. However, the framework does not comply with the basic constraint of biological neural networks - the division of neurons into excitatory and inhibitory units (Dale's law), and fails to account for a variety of recurrent and homeostatic currents, known to be important for the neural function. Here, we extend the theory of predictive coding with spikes and develop functional E-I networks that incorporates several important biophysical properties of cortical ensembles. Besides obeying Dale's law, our derivations also give a computational account of slow excitatory and inhibitory currents, as well as homeostatic currents such as spike-frequency adaptation. Our solution for the E-I network results in firing rate distributions, coefficient of variation and patterns of noise correlations similar to the ones observed in the cortex. We also address the scaling of the network that preserves the coding error and the total number of emitted spikes. We find that scaling the firing threshold with the natural logarithm of the number of neurons preserves the output of the network in terms of coding and computation.

1-050. Co-tuned, balanced excitation and inhibition in olfactory memory networks

Claire Meissner-Bernard CLAIRE.MEISSNER-BERNARD@FMI.CH
Friedemann Zenke FRIEDEMANN.ZENKE@FMI.CH
Rainer W. Friedrich RAINER.FRIEDRICH@FMI.CH

Friedrich Miescher Institute for Biomedical Research

The olfactory cortex shows structural features of an autoassociative memory network and plays a key role in olfactory memory retrieval. The zebrafish homolog of olfactory cortex (telencephalic area Dp) transiently enters a state of "precise synaptic balance" during the presentation of an odor (Rupprecht et al., 2018). This state is characterized by large synaptic conductances (relative to the resting conductance) and by co-tuning of excitatory and inhibitory synaptic conductances of individual neurons in odor space and in time. Other key features of Dp dynamics include low firing rates, odor-specific population activity and a dominance of recurrent inputs from Dp neurons relative to afferent inputs from neurons in the olfactory bulb (OB). We build a simplified, yet biologically plausible spiking neural network model of Dp using these key experimental observations as constraints. To achieve co-tuning of excitation and inhibition at the level of individual neurons, we introduce structured connectivity by increasing connection probabilities and/or strength among ensembles of excitatory and inhibitory neurons.

Ensembles are therefore structural memories of activity patterns representing specific odors. One key feature of olfactory processing is the retrieval of stored information from partial or degraded sensory inputs. Our preliminary results indicate that co-tuned excitation and inhibition improves pattern completion. Furthermore, attraction to stored patterns does not persist after removal of the odor in co-tuned networks. Hence, this study provides valuable insights into the computations performed by the olfactory cortex, as well as the effects of balanced state dynamics in associative memory networks.

1-051. Optimal small programs for computationally-constrained inferences that maximize behavioral utility

Tzuhsuan Ma[1] MAT@JANELIA.HHMI.ORG
Ann Hermundstad[2] HERMUNDSTADA@JANELIA.HHMI.ORG

[1]HHMI Janelia Research Campus
[2]Janelia Research Campus

To guide effective behavioral policies in the face of uncertainty and change, the brain is thought to make inferences about ambiguous surroundings to inform future actions. In such partially-observable settings, optimal strategies can be derived via two separate approaches: Bayesian inference for updating beliefs about unobservable properties of the environment, and value iteration for specifying optimal actions derived from those beliefs (Fig 1a). However, maintaining such internal beliefs requires infinite precision. Given limited computational resources, the brain must develop efficient strategies for determining not only how best to act, but whether, when, and how to devote resources to making inferences.

We formalize the interplay between limited resources and behavioral utility in partially-observable, nonstationary settings. We enumerate small programs that guide future actions based on past observations (Fig 1b) and identify optimal programs that maximize different utilities given a finite number of internal states M. We illustrate the value of this approach with a dynamic foraging task in which the type and probability of rewards—controlled by unobserved world states—changes over time (Fig 1c). To our surprise, the optimal M-state agent is not simply a discretized Bayesian agent; rather, it uses a fundamentally different program to achieve much higher utility in a two-port task (Fig 2).

We further demonstrate this difference in a task that requires tradeoffs between two action types (Fig 3). While a Bayesian agent uses the second action type as soon it delivers a slightly higher payoff, a small agent must balance the cost of reallocating resources and implementing relevant inferences, resulting in a strategy that requires a bigger payoff to warrant the second action. Together, this work shows how resource limitations necessitate a tight relationship between inference and action, and it informs the design of adaptive, modular agents that use small programs to exploit environmental structure.

1-052. Recurrent neural network model of cognitive control: effects of excitation-inhibition perturbations

Daniel Ehrlich[1,2] DANIEL.EHRLICH@YALE.EDU
John Murray[1] JOHN.MURRAY@YALE.EDU

[1]Yale University
[2]Interdepartmental Neuroscience Program

The ability to flexibly modulate behavioral responses to stimuli based on context provides one of the fundamental building blocks of cognition. How brains implement context-dependent cognitive control using fixed neural circuitry remains uncertain. Here, we trained an artificial recurrent neural network (RNN) on a well-studied and clinically relevant cognitive control task, the dot-pattern expectancy (DPX) task, and identified populations of units crucial to performance. In the DPX task the subject must use a cue to determine the response to a delayed probe. By disproportionately including trials of a single cue-probe pair the DPX task creates a prepotency bias, testing a subject's ability to exert cognitive control. Our trained RNN utilized populations of units that mirrored archetypes of neurons recorded in monkeys performing the DPX task, including switch units that flip their cue tuning between cue and probe epochs and units that maintain cue-tuned persistent activity through the delay. We perturbed this RNN via reduction in recurrent synaptic strengths, observing a pattern of behavioral errors that match both monkeys under NMDA-receptor antagonism and patients with schizophrenia. The perturbed network failed to gate out irrelevant probe information, and was diminished in switch units. Despite this, cue identity was decodable through the probe epoch, indicating no failure of working memory. Finally, we demonstrated the necessity of switch units to mediate gating by adding a rapid excitatory pulse to switch units during the probe epoch in the perturbed network. This treatment was sufficient to recover accurate behavior despite continued influence of the synaptic perturbation. Models of functional circuits provide a substrate to understand how synaptic-level perturbations such as changes in excitation-inhibition ratio may impact circuit computation and behavior. Here we demonstrate how a task-optimized circuit develops specific failure modes that match pharmacological and clinical deficits, and can be analytically explained and repaired in silico.

1-053. Neuronal refractoriness suppresses cortical variability in finite-size neural field model

Bastian Pietras[1,2] PIETRAS@TU-BERLIN.DE
Tilo Schwalger[1,2] TILO.SCHWALGER@BCCN-BERLIN.DE

[1]TU Berlin
[2]Institute for Mathematics

Fluctuations of neural population activity are widely observed in cortical recordings. Such cortical variability manifests through trial-by-trial variability and noise correlations, and critically affects the information capacity of neurons. Both experimental and theoretical studies aim at understanding the emergence as well as suppression of cortical variability, but the underlying neural mechanisms remain elusive. At present, it is poorly understood theoretically how internally generated variability depends on finite-size fluctuations and single neuron dynamics in spiking neural networks. Multi-attractor, or ring attractor, models have successfully reproduced experimental findings concerning, e.g., the diffusive wandering of bump activity in visual working memory tasks or the stimulus-dependent variance suppression in the visual cortex. But these approaches either use detailed computational models that are analytically intractable, or employ heuristic stochastic neural field models based on firing rates that lack a clear link to important network and neuron properties and drastically fail when trying to capture spike synchronization effects and nonstationary responses of the population activity. Here we derive an analytically tractable neural field model for finitely many, probabilistically spiking neurons. The reduced ring attractor model permits a comprehensive analysis including the stochastic wandering of bump activity as well as variance suppression upon stimulus onset as observed in experiments. Besides deriving a new neural field model valid for finite-size populations of spiking neurons, we demonstrate analytically how refractoriness—a fundamental property of neurons—contributes to the stimulus-induced suppression of cortical variability

1-054. A sampling-based model for optimal decision making

Camille Rullan[1,2] CER454@NYU.EDU
Cristina Savin[1,3] CS5360@NYU.EDU

[1]New York University
[2]Center for Neural Science
[3]CNS, CDS

Under a Bayesian framework, decision making involves computing with and marginalizing over complex, multidimensional posterior distributions that reflect the statistical regularities of the natural world. Sampling-based inference models provide greater flexibility in representing probability distributions than traditional parametric solutions. However, as they encode uncertainty implicitly, distributed across time and neurons, it remains unclear how such representations can be used for decision making. Here, we lay out a simple model for how a spiking neural network can use posterior samples to support Bayes-optimal decision making. We adopt the distributed sampling scheme proposed by Savin and Deneve (2014), in which marginalizing out nuisance variables reduces to appropriate linear projections of neural activity. In a second spiking circuit, we map samples of a task-specific marginal distribution into an instantaneous representation of uncertainty that can be readily used for decision making via a procedure inspired by kernel density estimation. We present a proof-of-concept network for a Gaussian posterior, and show that firing rates of neurons evolve in a manner that resembles what has been observed in evidence integration studies. We use this model to study how this computation is affected by stimulus dynamics. As the integration of posterior samples in the decision circuit is continuous in time, it leads to systematic biases after abrupt changes in the stimulus. This is reflected in behavioral biases towards recent history, similar to documented sequential effects in human decision making, and stimulus-specific neural transients. Overall, our work provides a first mechanistic model for decision making using sampling-based codes. It is also a stepping stone towards unifying sampling and parametric perspectives of Bayesian inference.

1-055. Neuronal mechanisms for sequential activation of memory items: Dynamics and reliability

Elif Koksal Ersoz[1,2] ELIF.KOKSAL-ERSOZ@INSERM.FR
Carlos Aguilar[3] CAGUILAR@MANTU.COM
Pascal Chossat[4] PASCAL.CHOSSAT@INRIA.FR
Martin Krupa[4] MACIEJ.KRUPA@INRIA.FR
Frederic Lavigne[5] FREDERIC.LAVIGNE@UNIV-COTEDAZUR.FR

[1]Inserm
[2]U1099
[3]Lab by MANTU
[4]INRIA Sophia Antipolis, Universite Cote d'Azur
[5]Universite Cote d'Azur

We present a biologically inspired model of activation of memory items in a sequence. Our model produces two types of sequences, corresponding to two different types of cerebral functions: activation of regular or irregular sequences. The switch between the two types of activation occurs through the modulation of biological parameters, without altering the connectivity matrix. Some of the parameters included in our model are neuronal gain, strength of inhibition, synaptic depression and noise. We investigate how these parameters enable the existence of sequences and influence the type of sequences observed. In particular we show that synaptic depression and noise drive the transitions from one memory item to the next and neuronal gain controls the switching between regular and irregular (random) activation.

1-056. Auto-associative memory without persistent activity in the olfactory system with co-tuned excitation and inhibition

Yue Kris Wu YUE.WU@FMI.CH
Rainer W. Friedrich RAINER.FRIEDRICH@FMI.CH
Friedemann Zenke FRIEDEMANN.ZENKE@FMI.CH

Friedrich Miescher Institute for Biomedical Research

The olfactory system is widely believed to implement an auto-associative memory network capable of generating similar representations to chemically distinct but functionally related odorants. Classic auto-associative memory models rely on strong recurrent excitatory connections accompanied by multi-stable attractor dynamics. In these models, selective activity persists even after removing the stimulus. However, activity patterns in memory-related neural circuits are often transient and can change rapidly. Moreover, experimental results revealed co-tuning of excitatory and inhibitory synaptic inputs to individual neurons, even in non-topographic memory networks such as the zebrafish homolog of olfactory cortex (Rupprecht and Friedrich, 2018). This specific organization of excitation and inhibition is usually not considered in models of memory networks. To explore transient dynamics and co-tuning in memory networks, we used a dynamical systems approach to analyze the response of supralinear recurrent networks to external stimuli. We found that powerful associative but transient responses occur when network dynamics switch from stable to unstable by additional afferent inputs but are re-stabilized through short-term synaptic plasticity. Spike-frequency adaptation, on the other hand, is insufficient to re-stabilize network dynamics. Next, we investigated the computational benefits of co-tuning for information processing. We demonstrated that in contrast to global inhibition, co-tuned inhibition broadens the parameter regime in which networks exhibit no persistent activity and forms robust representations regardless of the input strength. In summary, our analysis sheds light on the mechanistic effects of short-term synaptic plasticity and co-tuning for associative memory function in sensory networks and makes several testable predictions.

1-057. Distinct neural codes for attentional choice bias in visual and frontal cortex

Adithya Narayan Chandrasekaran[1,2] ADITHYANC@IISC.AC.IN
Ayesha Vermani[3] AYESHA.VERMANI@STONYBROOK.EDU
Nicholas Steinmetz[4] NSTEINME@UW.EDU
Tirin Moore[5] TIRIN@STANFORD.EDU
Devarajan Sridharan[1] SRIDHAR@IISC.AC.IN

[1]Indian Institute of Science, Bangalore
[2]Centre for Neuroscience
[3]Stony Brook University

[4]University of Washington, Seattle
[5]Stanford University

Attention enables efficient selection of task-relevant information. Yet, “attention” is not a unitary phenomenon. Behavioral studies indicate that attention operates through one of at least two mechanisms: By enhancing perceptual sensitivity (signal-to-noise) for task-relevant sensory information, or by enhancing choice bias (decisional gating) for task-relevant information, or both [1,2]. Neural substrates mediating these attention components remain controversial [3-7]. To address this issue, we analyzed neuronal activity in two cortical areas thought to contribute to sensitivity and bias, specifically, visual and prefrontal cortex. Neuronal activity was recorded simultaneously in visual (area V4) and prefrontal (frontal eye field, FEF) cortex in two monkeys performing a multialternative attention task. Animals covertly monitored four oriented gratings and reported the location of orientation change with an “anti-saccade” response. Attention was directed with a spatial probabilistic cue. Sensitivity and bias were quantified with signal detection theory [5], and found to be robustly dissociated across spatial locations. Choice bias was highest at the cued location (location of likely change), as expected. However, sensitivity was highest at the cue-opposite location (location of likely saccade). In addition, V4 neuronal activity modulated robustly with the magnitude of bias but, surprisingly, not with sensitivity at the cued location. FEF neuronal activity modulated with neither bias nor sensitivity; rather, it consistently signaled the location of the highest bias across space. In addition, trial-wise “noise” correlations between V4 and FEF neurons modulated strongly with bias, but not with sensitivity at the cued location. These results demonstrate that when the locus of spatial attention and saccade preparation are delinked, sensitivity and bias can be spatially dissociated. Contrary to recent literature [3], our results show that V4 neurons encode bias modulations also. Coordinating the flow of task-relevant information between the V4 and FEF may enable dynamic control of choice bias in the brain.

1-058. Temporal dynamics of visual statistical learning in the human hippocampus

Wei Tang[1,2] WT1@IU.EDU
Zhenghan Qi[3] ZQI@UDEL.EDU

[1]Indiana University-Bloomington
[2]Department of Computer Science
[3]University of Delaware

Humans are equipped with exquisite sensitivity to probabilistic information of sensory inputs. This ability – namely statistical learning (SL) – is evident across sensory modalities and different developmental stages [1,2]. Recent research suggests that the hippocampus plays a central role in SL, in that higher transitional probability (TP) between events corresponds with increased representational similarity in the hippocampal responses to these events [3,4,5]. Yet, the dynamics of the learning process, i.e. the temporal progression in which representational similarity builds up, remains unknown. In this study we investigated how the SL process unfolds by looking at the temporal structure of hippocampal activity. Blood-oxygenation-level-dependent (BOLD) data were acquired from 20 subjects performing a visual SL task. We developed a conditional-gaussian hidden Markov model (HMM) to decode TP from the BOLD signal. Critically, in the HMM, the emission probability of observed BOLD activity at each timepoint t was set to depend on the activity at timepoint t-1 as well as the covariance function of the two timepoints. Thus, the HMM decoding accuracy highly depended on how well the current state of BOLD activity was predicted by the previous state, and learning would presumably increase this predictability. We found that hippocampal voxels with significant decoding accuracy were clustered in the CA2/3/DG region. These voxels showed increasing predictability in their BOLD signals as a function of exposure to the structured stimuli. At the individual subject level, the autocorrelation of BOLD activity at timepoints that marked event boundaries was predicative of subjects' performance, such that greater autocorrelation corresponded with shorter reaction time. These results provide a detailed trajectory of SL in the CA2/3/DG region, opening the door to a novel HMM approach in studying SL dynamics in the human brain.

1-059. Development of reliable representations and their consistency with the intrinsic network structure

Sigrid Tragenap[1] TRAEGENAP@FIAS.UNI-FRANKFURT.DE
David E. Whitney[2] DE1WHITNEY@MSN.COM
Bettina Hein[3] HEIN@FIAS.UNI-FRANKFURT.DE
Jeremy T. Chang[3] JEREMY.CHANG@MPFI.ORG
David Fitzpatrick[4] DAVID.FITZPATRICK@MPFI.ORG
Matthias Kaschube[5] KASCHUBE@FIAS.UNI-FRANKFURT.DE

[1]Frankfurt Institute for Advanced Studies (FIAS)
[2]Max Planck Florida Institute, Jupiter, Florida, USA
[3]Neuroscience, Columbia University
[4]Max Planck Florida Institute of Neuroscience
[5]Frankfurt Institute for Advanced Studies

Experience plays an important role in cortical development, but does not operate on a blank slate. In animals with a modular network structure for orientation selectivity in visual cortex, patterns of activity from endogenous sources, including retina and LGN, support a robust modular structure, evident in spontaneous activity already several days prior to eye-opening [1]. How does visually evoked activity interact with this intrinsic network structure to produce reliable stimulus representations? To address this question, we employed chronic calcium imaging in ferret visual cortex around eye-opening [1,2]. Unlike the classical model that visually evoked activity in the early cortex is broad and poorly selective, we found that already several days prior to eye-opening grating evoked activity is as pronounced and modular as in the mature cortex. However, these early evoked patterns were highly variable for repeated presentations of the same stimulus. Moreover, they were only weakly similar to the intrinsic network structure, i.e. weakly overlapped with the prevalent variance components of spontaneous activity. In contrast, after several days of visual experience grating responses became highly consistent with the intrinsic network structure. This process was paralleled by a strong increase in reliability of evoked responses over the same period. A correlation between consistency and response reliability was evident even when comparing the responses to different stimuli at a given day. Finally, we found that activity patterns that were more similar to the intrinsic network structure were also more stable within a trial when compared to dissimilar patterns, possibly indicating a differential effect in driving circuit plasticity that strengthens the consistency between the evoked activity and the intrinsic networks. Thus, our results suggest a dynamic reorganization of cortical circuits, aligning stimulus evoked activity with the intrinsic network structure, a state associated with highly reliable grating evoked responses.

[1] Smith, G. B., Hein, B., Whitney, D. E., Fitzpatrick, D., & Kaschube, M. (2018). Distributed network interactions and their emergence in developing neocortex. Nature neuroscience, 21(11), 1600-1608 [2] Chang, J. T., Whitney, D., & Fitzpatrick, D. (2020). Experience-dependent reorganization drives development of a binocularly unified cortical representation of orientation. Neuron.

1-060. Olfactory evidence accumulation in mice

Hao Wu[1,2] HWU01@G.HARVARD.EDU
Siddharth Jayakumar[1] SJAYAKUMAR@G.HARVARD.EDU
Paul Masset[1,3] PAUL_MASSET@FAS.HARVARD.EDU
Joseph Zak[1] JZAK@FAS.HARVARD.EDU
Julien Grimaud[1] JGRIMAUD@G.HARVARD.EDU
Vikrant Kapoor[1] VIC.KAPOOR@GMAIL.COM
Venkatesh Murthy[1] VNMURTHY@FAS.HARVARD.EDU

[1]Harvard University
[2]Molecular and Cellular Biology
[3]Department of Molecular and Cellular Biology

In nature, odor cues from distant objects are sparse and highly fluctuating due to turbulent airflow. Animals may integrate odor concentration sampled over time rather than rely on transient odor concentration to effectively locate an object. To study how animals integrate and weigh discrete olfactory evidence over time, we developed a new behavioral task in which mice make binary decisions under fluctuating odor stimuli over many seconds. A custom-built device allowed us to precisely deliver discrete, short pulses of odors at arbitrary Poisson-distributed pulse rates. We found that trained mice can readily differentiate stochastic odor stimuli with different average pulse rates presented over many seconds. In order to investigate how active, discrete sniff-based sampling of a stochastically-varying environmental cue affects the neural representation and perceptual interpretation of the cue, we performed calcium imaging in the axon terminals of olfactory sensory neurons (OSNs) in the glomeruli of olfactory bulb (OB). We discovered that OSN activity was highly-modulated by the phase of the sniffing cycle.

Regression of behavioral outcome against the timing of odor pulses in the breathing cycle revealed a kernel that weighted pulses arriving during the inhalation cycle more than during exhalation. This kernel matched the OSN activity kernel over breathing cycle, suggesting that the strength of the perception elicited by single pulses was directly related to the strength of the OSN responses. Finally, decision noise scaled with the number of pulses presented. Our study indicates that mice integrate discrete olfactory inputs over several seconds to make decisions and that perceptual evidence is weighted by the intensity of the OSN response to the input. Furthermore, our platform introduces a new paradigm in perceptual decision-making in which we can, unlike in vision or audition, record neural activity at all levels, from the first layer of sensory neurons to the decision-making networks.

1-061. Distinguishing neural signatures of population codes encoding uncertain information

Qiao Zheng[1] QIAO_ZHENG@HMS.HARVARD.EDU
Jeff Beck[2] BAYESIAN.EMPIRIMANCER@GMAIL.COM
Jan Drugowitsch[1,3] JAN_DRUGOWITSCH@HMS.HARVARD.EDU

[1]Harvard Medical School
[2]Duke University
[3]Neurobiology

Ample behavioral evidence indicates that our brains perform probabilistic computations, but there is little agreement regarding underlying neural representation. Probabilistic population codes (PPCs), distributed distributional codes (DDCs) and quantile/expectile codes all suggest that probability distributions are represented by their parameters or statistics, but differ in which of these they encode. Past theoretical work has identified neural activity features characterizing PPCs, but similar characterization is scarce for other codes, making it hard to distinguish them in neural data. We address this gap by identifying fundamental neural activity features characterizing the codes. We first show analytically that a linear change in population activity results in a linear change in log-probabilities in PPCs, in probabilities in DDCs, and in an overall shift of the distribution in quantile/expectile codes. Furthermore, an increase in uncertainty (i.e. entropy) is reflected by lower population activity in PPCs, while in DDCs this is not the case in general. Lastly, while the decoding model underlying PPCs is compatible with common neural encoding models (e.g., independent Poisson), DDCs do not feature this compatibility.

To explore how these features could make neural codes distinguishable in experiments, we simulated neural activity encoding known probability distributions. Model comparison correctly identified the types of the generating codes, and revealed that such identification is only possible if the set of encoded distributions is sufficiently diverse. Applying this approach to neural population activity recorded from mouse V1 in response to drifting gratings, we could not identify the code type as the encoded ground truth distributions is unknown, but nonetheless found that a PPC-type code was able to recover the drift direction information best. Overall, we expect our work to support and inform how to identify the representations underlying uncertain information in the brain, and how they can be used for computations to feature efficient behavior.

1-062. Coding of anxiety states by coordinated neural ensembles in the mPFC

Garima Shah[1,2] GSHAH5@JHU.EDU
Juliana Popovitz[3] POPOVITZ@GMAIL.COM
Barbara Holt[4] BHOLT4@JHU.EDU
Shreesh Mysore[4] MYSORE@JHU.EDU
Hita Adwanikar[4] HITA@JHU.EDU

[1]JOHNS HOPKINS UNIVERSITY
[2]Psychological and Brain Sciences
[3]Charles River Laboratories
[4]Johns Hopkins University

Distributed neural circuits in the prefrontal and limbic regions are known to be involved in the control of anxiety-like behaviors. Much less is known about the neural encoding of anxiety states, with the primary finding, thus far, being the presence of individual neurons that are selective for high or low anxiety states in the medial prefrontal cortex, the amygdala and the ventral hippocampus. The complexity of an affective state such as anxiety suggests that additional insights may be gained from the investigation of coding by ensemble activity patterns - an approach that has proved effective in studies of other encoding problems in neuroscience. Inspired by these, here, we examine the joint neural ensemble activity patterns (or 'code') in mPFC using calcium imaging in mice engaged in a classic assay of anxiety, the elevated zero maze. Using a miniature fluorescence microscope, we longitudinally tracked the Ca^{2+} dynamics of mPFC neurons in freely moving mice as they traversed the exploration or avoidance zones

of the elevated zero maze (EZM). Our data reproduce the classic results of subsets individual neurons being selective but go beyond to show that the joint activity patterns are distinct. These patterns occupy distinct subsets of n-d space, with the open arm representational cloud being denser (more stereotypical patterns), than that of the closed arm. An additional powerful finding is that the distinctive coordinated activity patterns between the two arms is not purely driven by the selective neurons. Surprisingly, sub-ensembles of neurons that are not individually selective also convey distinctive information about anxiety states. Results reveal a novel, distributed, coordinated neural code for anxiety states in the mPFC.

1-063. Real-time prediction of multifunctionality in Aplysia californica feeding using a Boolean model framework

Victoria Webster-Wood[1,2] VWEBSTER@ANDREW.CMU.EDU
Jeffrey Gill[3] JEFF.GILL@CASE.EDU
Peter Thomas[3] PJTHOMAS@CASE.EDU
Hillel Chiel[3] HJC@CASE.EDU

[1]Carnegie Mellon University
[2]Mechanical Engineering
[3]Case Western Reserve University

The ability to continuously switch and adapt behaviors is critical for animals to robustly operate in complex environments and requires coordination between circuits in the central nervous system and biomechanical properties of the body. Such adaptable, multifunctional behavior is observed in Aplysia californica feeding. Aplysia exhibit three unique classes of behaviors: (1) biting – an unsuccessful attempt to grasp food, (2) swallowing – successfully grasping and ingesting food, and (3) rejection – moving inedible food out of the feeding apparatus. Additionally, swallowing behaviors adapt to mechanical loading. Aplysia feeding is a tractable model for studying multifunctional neuromuscular control. Computational models of the nervous system are a common tool for testing hypotheses in highly controlled environments. However, many existing modeling approaches are computationally expensive, and, as a consequence, cannot be tuned or simulated in real-time to fit individual animals. Furthermore, they are not, in general, connected to the biomechanics of the body and thus cannot predict detailed behavioral outputs. To predict and control morphologically complex peripheries and multifunctional behaviors, computationally efficient models of neuromuscular systems are needed. Towards this goal, we have developed a hybrid Boolean/continuous modeling framework capable of representing neural bursting activity and associated biomechanics substantially faster than real-time. Using this framework, we have implemented key neural circuitry and musculature involved in Aplysia feeding coupled to a simple biomechanical model. The model is capable of qualitatively reproducing the three multifunctional feeding behaviors (biting, swallowing, and rejection), switching between behaviors in response to external sensory cues, and adjusting behaviors based on external mechanical loading. Model runtimes are 2-3 orders of magnitude faster than real-time on standard hardware. This framework is computationally efficient while incorporating known neural circuitry and feeding apparatus biomechanics. The model is easily expandable and may provide a future foundation for real-time prediction and control of animal experiments.

1-064. Cognitive strategies for resolving uncertainty in the presence of hierarchies

Mahdi Ramadan[1,2] MRAMADAN@MIT.EDU
Mehrdad Jazayeri[1] MJAZ@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Brain and Cognitive Sciences

Converging evidence suggests that humans and monkeys can revise their decisions when sensory evidence contradicts their beliefs. However, the vast majority of work on this topic comes from tasks with relatively simple sensorimotor contingencies. In contrast, we know very little about the reasoning strategies that are at play when facing more complex decision trees with multiple levels of hierarchy. Previous experiments tackling this question used tasks that were either too simple to distinguish between candidate models, or too complex to accommodate quantitative modeling. Here, we tackle this problem using a novel inference task in which we can precisely control the uncertainty at different levels of hierarchy. Monkeys had to infer the path of an invisible ball within a H-shaped maze based on a sequence of visual cues that signaled when the ball changed direction. The maze was structured hierarchically so that monkeys could use each visual cue to validate or revise their belief about the ball position. We examined behavioral responses in relation to two cognitive models that were broadly consistent with the behavior: 1) a retrodictive model that uses later segments of the maze to retroactively resolve earlier

uncertainty, and 2) a model that uses counterfactual reasoning to consider alternatives. General comparison of the monkeys' behavior to model simulations provided qualitative evidence in support of the counterfactual model. Further analysis of the behavior using model fitting and comparison provided further quantitative evidence that monkeys relied on a counterfactual reasoning strategy to revise their decisions. Together, these results provide a simple yet expressive example of how to combine careful task design with detailed modeling to distinguish between high-level cognitive reasoning strategies.

1-065. Learning dendritic opinion weighting by cortical neurons

Jakob Jordan[1,2] JAKOB.JORDAN@UNIBE.CH
Joao Sacramento[3,4] SACRAMENTO@INI.ETHZ.CH
Mihai Petrovici[5,6] MIHAI.PETROVICI@UNIBE.CH
Walter Senn[7] WALTER.SENN@UNIBE.CH

[1]University of Bern
[2]Department of Physiology
[3]University of Zurich & ETH Zurich
[4]Institute of Neuroinformatics
[5]University of Bern and Heidelberg University
[6]Department of Physiology and Kirchhoff-Institute for Physics
[7]Universitat Bern

Successful decision making involves integrating a variety of different opinions, reliable to varying degrees. How should one combine these to maximize the chances for a desired outcome? One option is to weight each opinion according to its estimated reliability.

We suggest that such opinion weighting is naturally performed by cortical neurons. Each dendritic branch forms a local membrane voltage, the analog of an opinion. The reliability of a dendritic branch is encoded in the local membrane conductance. The biophysics of the bidirectional voltage propagation then leads to the formation of a somatic membrane voltage representing a weighted combination of dendritic opinions. A neuron's output consequently reflects decisions based on the combined dendritic opinions.

In a probabilistic framework we derive somatic membrane potential dynamics which sample from a product model of dendritic probability distributions. For the case of Gaussian noise, the average somatic membrane potential is identical to a Bayes-optimal maximum-a-posteriori estimate (e.g., Knill & Pouget, 2004). We derive a local synaptic plasticity rule which allows the somatic potential to approximate specific target distributions via gradient descent, and additionally assigns appropriate relative reliabilities to different cortical pathways converging on a neuron (cf. Friston, 2018).

We demonstrate successful learning of a prototypical opinion weighting task, namely the integration of multimodal sensory cues to guide behavior. The trained model provides normative interpretations of behavioral and neuronal data from cue integration experiments and makes specific predictions about system behavior and single cell dynamics. The conductance-based nature of synaptic coupling hence may not be an artifact of the biological substrate, but rather enable neurons to perform computations previously thought to be realized at the circuit level. Furthermore, the mathematical and computational tractability of our framework allows its extension to hierarchical and recurrent implementations of neuronal opinion weighting.

1-066. Brain-inspired architectures for efficient and meaningful learning from temporally smooth data

Shima Rahimi Moghaddam[1,2] SHIMA@JHU.EDU
Fanjun Bu[1] FBU2@JHU.EDU
Christopher Honey[1] CHRIS.HONEY@JHU.EDU

[1]Johns Hopkins University
[2]Psychological and Brain Sciences

In the human brain, internal states are autocorrelated over time, due to local recurrence and other intrinsic circuit properties. At first glance, temporal smoothness of internal states presents a problem for learning input-output mappings (e.g. image categorization), because a learning system must associate an output (e.g. the target label for image n) with a mixture of current and prior inputs (e.g. images n and n-1). How does such temporal "smoothness" of internal states affect the efficiency of learning? How does it affect the kinds of representations that are learned? We found that neural networks whose internal states are mixed over time (via "leaky memory" of internal states) can learn more efficiently than memoryless networks, when they are trained using input data that is temporally “smooth”, like data in in the real world. Moreover, we found that when a memory-resetting mechanism was added to leaky memory networks, they could flexibly adapt to the temporal properties of the training data. Finally, we found that networks with multi-scale leaky memory and memory- gating could learn internal representations that “un-mixed” data sources which varied on fast and slow timescales. Altogether, we showed how brain-inspired mechanisms enable neural networks to learn more efficiently from temporally smooth data, and to generate internal representations that separated timescales in the training signal.

1-067. Reward resonance in amygdala-based circuits

Jeremiah Morrow[1,2] MORROWJE@OHSU.EDU
Vincent D. Costa[1] COSTAV@OHSU.EDU

[1]Oregon Health and Science University
[2]Behavioral Neuroscience

The ability to learn from trial and error feedback is dependent on the basolateral amygdala (BLA), ventral striatum (VS), and orbitofrontal cortex (OFC) but how information is communicated between populations of neurons in these interconnected regions is not well understood. Empirical and theoretical work has demonstrated that interactions between low and high frequency oscillations in local field potentials (LFPs), known as cross frequency coupling, may serve as a mechanism to facilitate information transfer between interconnected neuronal networks by linking system wide inputs and outputs. In particular, coupling between the phase of theta oscillations and the power of gamma oscillations, a phenomenon known as phase-amplitude coupling (PAC), has been linked to a variety of biological functions. Here, we recorded LFPs from the BLA, VS, and OFC while monkeys performed a three-armed bandit task to determine whether PAC within and between these structures is modulated as a function of learning. Specifically, we examined how PAC values were modulated as a function of reward prediction errors (RPEs), a key reinforcement learning signal. We found that coupling of theta and gamma activity within and between the BLA and VS or OFC was robustly modulated by RPEs during the intertrial interval between successive choices. We found a negative correlation between PAC values and RPE. Altogether, these results suggest that reward information and value updating synchronizes neural networks operating at distinct frequency bands to facilitate information transfer between key nodes of the motivational brain.

1-068. Robust and flexible cortical responses to single spike activations

Juan Luis Riquelme[1] JUAN.RIQUELME@BRAIN.MPG.DE
Mike Hemberger[1] MIKE.HEMBERGER@BRAIN.MPG.DE
Gilles Laurent[1] GILLES.LAURENT@BRAIN.MPG.DE
Julijana Gjorgjieva[1,2] GJORGJIEVA@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]Computation in Neural Circuits Group

Experiments in the mammalian cortex have shown that a single spike may increase circuit firing rates. In the reptilian cortex ex vivo, a spike can trigger repeatable sequences of neuronal activations in the surrounding network. However, the mechanisms underlying this reliable propagation are not fully understood, and the computational role of these firing sequences is not clear.

We explored spike propagation in simulations of large networks fitted to reptilian experimental data. Our networks included rare but powerful synapses from a long-tailed distribution. The model reproduces the single-spike-evoked sequences observed experimentally under low spontaneous activity. Our model can also generate sequences at higher activity levels. However, inhibitory neurons are particularly sensitive to high activity and become less reliable. We thus predict that spike sequences likely propagate among excitatory but not inhibitory neurons in vivo. Examining the activation paths, we observed that interneuron spikes result mainly from many weak converging inputs, whereas pyramidal neuron activations rest on rare but strong connections. Simulating networks with truncated synaptic-strength distributions, we found that sparse networks of strong connections are sufficient to generate sequences. By contrast, dense networks of weak connections increase overall activity, reducing reliability.

We propose that the opposing roles of strong and weak connections may form a basis for the routing of activity. Indeed, when we clustered sequentially activated neurons by their activity, we identified, within each sequence, neuron groups that could be activated separately. Examining the paths of excitation within these groups revealed branches of strong connections. These branches are weakly and sparsely connected with one another, enabling the selective gating of activity propagation between them. Our work details how a single spike can robustly trigger a response in a network and how that response can be flexibly gated, providing a mechanism for routing information in cortical circuits.

1-069. How a mouse licks a spout: Cortex-dependent corrections as the tongue reaches for, and misses, targets

Brendan Ito[1,2] BSI8@CORNELL.EDU
Teja Bollu[3] TBOLLU@SALK.EDU
Brian Kardon[1] BMK27@CORNELL.EDU
Jesse Goldberg[1] JHG285@CORNELL.EDU

[1]Cornell University
[2]Neurobiology and Behavior
[3]Salk Institute for Biological Studies

Accurate goal-directed behavior requires the constant monitoring and correction of ongoing movements. For example, when reaching to uncertain or displaced targets, errors are estimated and compensated for in real time, resulting in corrective submovements (CSMs) that redirect the hand to the target. In primate reach tasks, CSMs occur in conditions where corrections can be planned in advance – such as when a target is uncertain or far away – and also in conditions when CSMs are generated on-the-fly, such as in 'double-step' experiments when a target unexpectedly jumps mid-reach. To examine the role of cortex in online corrections, we adapted the 'double-step' paradigm to a lick task for mice. To do this, we combined kilohertz framerate imaging and a deep neural network to track tongue kinematics in 3D with decamicron-millisecond spatiotemporal precision. While tracking the tongue, we detected the offset of tongue-spout contact on the first lick (L1) within a lick bout in real time and rapidly retracted the spout so that by the onset of the second lick (L2) the spout was 1 mm farther away. After tongue protrusions missed the spout on double-step trials, mice generated on-the-fly, within-lick CSMs that extended to reach the more distant spout. Photoinhibition of anterolateral motor cortex (ALM) impaired these corrections CSMs, and electrophysiological recordings in ALM revealed error and correction-related signals with appropriate timing to influence ongoing licks. Though less than a tenth of a second in duration, a single mouse lick exhibits hallmarks of online motor control associated with a primate reach, including cortex-dependent corrections after misses.

1-070. Feedback representations shape dynamics and performance in trained recurrent neural networks

Kabir Dabholkar[1] KABIR.DABHOLKAR@STUDENTS.IISERPUNE.AC.IN
Omri Barak[2] OMRI.BARAK@GMAIL.COM

[1]Indian Institute of Science Education and Research, Pune
[2]Technion, Haifa

Recurrent Neural Networks (RNNs) trained on neuroscience tasks are promising models of population dynamics of their biological counterparts. In these frameworks RNNs are only constrained by task definition, and can potentially find many solutions to the same task – differing in terms of both their neural signatures, and their dynamical properties (e.g. stability). This poses a problem when comparing them to experimental data – how can the RNN be expected to find a solution matching the biological one? Recent studies approached this problem by constraining part of the network to match data [1] or assuming certain regularizations [2]. The effect of these suggestions on network dynamics, however, is unknown. Here, we study families of networks that solve the same task (identical input-output functionality), but have different internal representations. We then manipulate these representations systematically and characterise their dynamical consequences. We construct these families within the reservoir computing paradigm, in which RNNs have random and structured parts. The structured part is a self-consistent, low-dimensional feedback loop for a signal – usually the task defined target output. Here, instead of the original output, we train this loop with modified signals while still requiring that a second readout (without feedback) produces the original output from the population activity. The result is a family of networks with identical functionality but (1) perturbation of varying ranks on the initial connectivity and (2) varying signals passing through the feedback loop. We analyze two tasks: sine-wave generation and multiple fixed-point formation. In both cases, we find a qualitative change in performance, noise robustness and local dynamics around the fixed points as we modulate the dimensionality of the feedback loop. Our findings suggests that analysing and modifying feedback signals can be a useful method to constrain RNN models to biological activity and can complement existing training frameworks.

1-071. Learning sensory representations for flexible computation with recurrent circuits

Liang Zhou LIANG.ZHOU.18@UCL.AC.UK
Jorge Menendez JORGE.MENENDEZ.15@UCL.AC.UK
Peter Latham PEL@GATSBY.UCL.AC.UK

Gatsby Computational Neuroscience Unit, UCL

A common approach to modeling time-dependent tasks is to optimize the connectivity of a recurrent neural network (RNN) so that it reproduces a desired behavior. But neural circuits are presumably optimized over a lifetime of experience to solve a wide range of ethologically relevant tasks. One might thus expect that, when learning a new task, this prior structure would be exploited for computation. This seems to be the case in motor cortical circuits, which are able to learn a new motor task by recycling the same activity patterns used during a previously learned task [Golub et al '18]. Here we propose a model in which a general-purpose RNN is exploited for solving new tasks, without needing to continuously re-optimize the recurrent connectivity. Instead, learning a new task boils down to learning the appropriate representation of the task-relevant stimuli that will drive the RNN to produce the correct responses to task cues. Moreover, incorporating low-dimensional structure in the stimulus can drastically reduce the number of parameters that need to be learned, leading to much faster learning. In particular, rather than having to learn all N^2 weights of the recurrent connectivity, only the weights specifying how the stimulus is represented and how the RNN's activity is read out need to be learned. This setting in principle lends itself well to biologically plausible learning algorithms. We empirically show that such an architecture can be utilized to accurately model behavioral performance on the ready-set-go task [Jazayeri & Shadlen '15], as well as a custom delay-copy task, with performance comparable to that of fully-trained RNNs.

1-072. Abstract number sense in untrained deep neural networks

Jaeson Jang[1,2] JAESONJANG@KAIST.AC.KR
Gwangsu Kim[1] KGSPIANO@KAIST.AC.KR
Seungdae Baek[1,2] SEUNG7649@KAIST.AC.KR
Song Min[1] NIGHT@KAIST.AC.KR
Se-Bum Paik[1] SBPAIK@KAIST.AC.KR

[1]Korea Advanced Institute of Science and Technology
[2]Department of Bio and Brain Engineering

Number sense, the ability to estimate numbers without counting, is observed in naive animals in the absence of learning but how this innate cognitive function emerges in the brain remains unclear. In monkeys and crows, neurons in the prefrontal cortex and other brain areas were observed to respond selectively to the abstract number of visual items ("numerosity"), suggesting that number-selective neurons ("number neurons") provide a foundation for an innate number sense in the brain. However, details of how this neuronal number tuning emerges before visual training are not yet understood. Here, we show that number selective neurons can arise spontaneously even in completely untrained deep neural networks, and that these neurons enable the network to discriminate numerosity. Using an AlexNet, an artificial deep neural network that models the ventral visual stream of the brain, we found that number selective neurons emerge in randomly initialized deep neural networks, and that they show the single- and multi-neuron characteristics observed in biological data following the Weber-Fechner Law. We found that these neurons can induce the abstract number sense, the ability to discriminate numerosity independent of low-level visual cues, and confirmed that the responses of these neurons enable the network to perform an abstract numerosity comparison task, even under the condition that the numerosity in the stimulus is incongruent with low-level visual cues. We found that the observed number tuning in a randomly initialized network originated from a combination of monotonically decreasing and increasing neuronal activities in the earlier layers, which emerges spontaneously from bottom-up projections. This implies that the number tuning emerges from the statistical variation of feedforward projections in hierarchical neural networks. Our findings provide novel insight into the origin of innate cognitive functions in the absence of learning.

1-073. Normative decision asymmetries with symmetric priors but asymmetric evidence

Tahra Eissa[1] TAHRA.EISSA@COLORADO.EDU
Joshua Gold[2] JIGOLD@PENNMEDICINE.UPENN.EDU
Kresimir Josic[3] JOSIC@MATH.UH.EDU
Zachary Kilpatrick[1,4] ZPKILPAT@COLORADO.EDU

[1]University of Colorado Boulder
[2]University of Pennsylvania
[3]University of Houston
[4]Applied Mathematics

Decision asymmetries can reflect many factors but are typically considered normative only when they result from asymmetric priors or values. Previous studies of normative decision asymmetries have focused on restricted sets of conditions that have limited our understanding of potential sources of asymmetry. Most studies assumed that available evidence is symmetric and evidence is matched in strength for each alternative. Here we examine decisions between two alternatives that are each associated with relatively rare, asymmetric evidence (e.g., deciding which of two slot machines gives higher odds of winning by watching their outputs). We show that when evidence is asymmetric and sparse, decision asymmetries are inevitable. With symmetric priors, these evidence-driven asymmetries can raise a conundrum for normative decision-makers, who must reconcile decision asymmetries with an expectation that no asymmetries should exist. We examined how 200 human participants handled this conundrum and compared their performance to normative and non-normative models.

We used a ball-drawing task in which balls of one color could be rare. Subjects saw a short string of balls drawn with replacement from one of two equally likely jars with known ratios of ball colors and were asked which jar was used. Most participants reported choices supported by more extreme (i.e., rarer) evidence more often, despite both alternatives being equally probable. This discrepancy was consistent with a Bayesian ideal observer that displays similar asymmetries in its decisions.

Notably, many subjects' asymmetries were enhanced relative to the ideal observer, matching Bayesian models that under-weighed the value of rare balls. Model-free results confirmed that deviations from optimality correlated with a decreased use of task-relevant information. These results provide quantitative and theoretically grounded insights into how humans use rare events to make inferences, relevant to predictions in both real-world situations (stock-market crashes) and common laboratory tasks (changes in reward contingencies).

1-074. Neuromodulatory control of spontaneous activity and circuit maturation in development

Jan Kirchner[1,2] JAN.KIRCHNER@BRAIN.MPG.DE
Paloma Maldonado[3] P.MALDONADO@NIN.KNAW.NL
Alvaro Nuno-Perez[4] ALVARO.NUNOPEREZ@UNIL.CH
Christian Lohmann[3] C.LOHMANN@NIN.KNAW.NL
Julijana Gjorgjieva[1,5] GJORGJIEVA@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]Computation in Neural Circuits
[3]Netherlands Institute for Neuroscience
[4]The University of Lausanne
[5]Computation in Neural Circuits Group

The sensory cortex of a newborn animal goes through an intricate process of circuit refinement. This process occurs mostly in the absence of sensory input and is instead guided by spontaneously generated neural activity, which shapes the circuit through activity-dependent plasticity. One prominent factor that can regulate neural activity is the neuromodulator oxytocin. Since oxytocin is prominently expressed in the developing sensory cortex, it is a compelling question how oxytocin might affect spontaneous activity and therefore, through activity-dependent plasticity, circuit maturation.

To investigate how oxytocin modulates activity in the developing cortex, recent in vivo calcium imaging of excitatory neurons in the visual cortex of neonatal mice found that exogenous application of oxytocin rapidly reduces the frequency of spontaneous activity. Furthermore, oxytocin reduces correlations between neurons in a distance-dependent manner, where distant pairs of neurons experience stronger reduction than nearby pairs. Subsequent patch-clamp recordings of inhibitory interneurons revealed that oxytocin increases excitability of the somatostatin-expressing (SST+) subpopulation. To understand how this subpopulation might control excitatory activity, we developed a recurrent spiking neural network model with excitatory and inhibitory populations, distance-dependent connectivity, and parameters based on experiments. We found that the increase in the resting potential of a small inhibitory subpopulation can explain the reduction in frequency of spontaneous activity as well as the distance-dependent reduction of correlations. We identified a local inhibitory connectivity motif responsible for these population-level changes. This motif allowed us to analytically compute correlations between excitatory neurons under changing inhibition. By implementing Hebbian plasticity between excitatory neurons, we further explored the possibility that oxytocin-induced reduction in correlation can instruct the refinement of lateral connectivity within the cortex. This data-driven modeling approach allows us to understand how neuromodulation of an inhibitory subpopulation affects network activity and allows us to make predictions about circuit connectivity when oxytocin signaling is perturbed.

1-075. Flexible use of memory by food-caching birds in a laboratory behavioral paradigm

Marissa Applegate[1,2] MCA2179@COLUMBIA.EDU
Dmitriy Aronov[3] DA2006@COLUMBIA.EDU

[1]Columbia University
[2]Neurobiology and Behavior
[3]Zuckerman Mind Brain Behavior Institute, Columbia University

The neural mechanisms of episodic memory, or memory of specific events, are challenging to study in laboratory conditions. One promising approach is to take advantage a behavior in which episodic-like memories are naturally used by the animal. Food-caching birds like black-capped chickadees cache up to thousands of food items per day and use memory to find their hidden caches. To utilize this behavior in the lab, we developed a food-caching arena for chickadees and modeled their use of memory in this arena. Our behavioral paradigm 1) allows birds to express their natural caching instincts, 2) permits high-throughput, automated behavioral tracking, and 3) is compatible with modern neuroscience techniques, such as awake-tethered recordings.

We recorded chickadee behavior in two tasks: one in which they freely cached seeds throughout the arena and one in which they retrieved all previously made caches. We then built maximum-likelihood models to explain the birds' behavior in these tasks using memory-based and non-memory-based parameters. We found that in the caching task, birds avoided caching into sites that were already occupied by previous caches. This effect was driven both by avoidance of known-occupied sites, and by attraction to known-empty sites, indicating a content-specific memory. In the retrieval task, chickadees instead preferentially returned to the occupied sites, showing that they use memory flexibly depending on task context. These results reveal a surprisingly flexible use of memory by chickadees and set up a powerful episodic memory paradigm for investigating neural mechanisms.

1-076. WavelAt: Mechanisms for the spatiotemporal modulation of visual resolution by exogenous attention

Alessio Franci[1,2] AFRANCI@CIENCIAS.UNAM.MX
Martina Poletti[3] MARTINA_POLETTI@URMC.ROCHESTER.EDU

[1]National Autonomous University of Mexico
[2]Mathematics
[3]University of Rochester

Recent experiments have revealed how exogenous attention modulates in time and space fine spatial vision across the central fovea. At short delays, visual resolution is locally enhanced at the cued location and decreased at un-cued locations. At long delays, the pattern is reversed. The first contribution of our work is to explain mechanistically how vision of detail can be controlled in this highly localized fashion. Our modeling hypothesis is the existence of tunable horizontal connections in V1, whose E/I-balanced gain-control is mediated by modulatory inputs. Frequency-domain methods show that a balanced scaling of E/I horizontal gains in V1 can focus (enhance the resolution of) or defocus (decrease the resolution of) the feedforward visual kernel originating in the upstream visual pathway, accordingly to a continuous wavelet zoom. We illustrate our theory on a neural field model of V1. The second contribution is to show that the existence of an attention traveling wave (WavelAt), originating at the cued location and coordinating the E/I-balanced modulation of V1 horizontal connections, explains and reproduces the spatiotemporal dynamics of visual-resolution control. Lateral inhibition around the bump of the WavelAt explains the early resolution decrease at un-cued locations. The after-hyperpolarization period of the WavelAt explains the late resolution decrease at the cued location. We briefly discuss the possible physiological origins of the WavelAt. Our modeling results constitute a novel mechanistic account of covert-attention-related resolution changes in the visual pathway and reveal an overlooked connection between visual resolution control and wavelet theory. Besides explaining a number of recent experimental observations, they allow us to formulate

novel predictions about the neuronal mechanisms behind spatiotemporal visual-resolution control, which we plan to test with further experiments.

1-077. Learning is shaped by an abrupt change in "neural engagement"

Jay Hennig[1,2] JHENNIG@CMU.EDU
Emily Oby[3] EMILYOBY@GMAIL.COM
Matthew Golub[4] GOLUB.MATTHEW@GMAIL.COM
Lindsay Bahureksa[1] LBAHUREK@ANDREW.CMU.EDU
Patrick Sadtler[3] PATRICK.T.SADTLER@GMAIL.COM
Kristin Quick[3] KRISTIN.QUICK@PITT.EDU
Stephen Ryu[4] SEOULMANMD@GMAIL.COM
Elizabeth Tyler-Kabara[3] ELIZABETH.TYLER-KABARA@CHP.EDU
Aaron Batista[3] AARON.BATISTA@PITT.EDU
Steven Chase[1] SCHASE@ANDREW.CMU.EDU
Byron Yu[1] BYRONYU@CMU.EDU

[1]Carnegie Mellon University
[2]Center for the Neural Basis of Cognition
[3]University of Pittsburgh
[4]Stanford University

Internal states such as arousal, attention, and motivation are known to modulate brain-wide neural activity, but how these processes interact with learning is not well understood. During learning, the brain must modify the neural activity it produces to improve behavioral performance. How do internal states affect the evolution of this learning process? To answer this question, we trained macaque monkeys to perform a brain-computer interface (BCI) task. Monkeys volitionally modulated neural activity, recorded in primary motor cortex (M1), to move a computer cursor to hit visual targets. In each experimental session, monkeys first used an "intuitive" BCI mapping, during which they demonstrated proficient cursor control. To induce learning, we then introduced a "perturbed" BCI mapping, which required monkeys to modify their neural strategy in order to regain proficient control. We identified large, abrupt fluctuations in neural population activity that tracked the monkey's pupil size, a correlate of internal state. These fluctuations drove population activity along dimensions we term neural engagement axes. In a BCI, the causal relationship between neural activity and behavior is known. This allowed us to understand how changes in neural engagement during learning impacted behavioral performance for different task goals. We found that neural engagement increased abruptly at the start of learning, and then gradually retreated, even when doing so negatively impacted behavioral performance. Surprisingly, when the abrupt increase in neural engagement led to immediate performance improvements, monkeys were able to maintain that performance increase even as neural engagement subsided on subsequent trials. Monkeys thus learned some task goals more quickly than others, in a manner that was predictable based on how changes in neural engagement impacted behavior (i.e., cursor speeds). These results suggest that changes in internal states, even those seemingly unrelated to goal-seeking behavior, can influence how behavior improves with learning.

1-078. Parieto-collicular interactions during perceptual decisions

Gabriel Stine[1,2] GMS2191@COLUMBIA.EDU
Eric M Trautmann[1] EMT2177@COLUMBIA.EDU
Michael N Shadlen[1] SHADLEN@COLUMBIA.EDU

[1]Columbia University
[2]Neuroscience

Many decisions arise through an accumulation of evidence to a terminating threshold. Previous studies showed that the lateral intraparietal area (LIP) represents the accumulation of evidence when monkeys make decisions about the direction of random dot motion (RDM) and express their decision with a saccade. However, LIP is interconnected with other brain regions that also display decision-related activity. The superior colliculus (SC) receives direct input from LIP and projects disynaptically to LIP via the pulvinar. Like LIP, many SC neurons display spatially selective, persistent activity and decision-related activity. We wished to ascertain whether SC conveys signals like those in LIP to the oculomotor system or if it transforms these signals in some way. We used multi-contact linear probes to record simultaneously from populations of single neurons with overlapping response fields in LIP and SC, as a monkey performed a reaction-time, RDM task. Unsurprisingly, LIP neurons reflected the accumulation of noisy motion-evidence. In contrast, SC neurons—including those with persistent activity (e.g., prelude neurons)—did not reflect evidence accumulation. They generated an all-or-none burst of activity immediately preceding the decision's report. Additionally, the saccadic burst was sometimes preceded by smaller,

premature bursts of activity, which resembled the dynamics of the saccadic burst but with lower amplitude. Both premature and saccadic bursts were associated with a sharp increase in LIP activity, but saccadic bursts were also preceded by an elevated baseline of LIP activity. We conclude that SC and LIP play fundamentally different roles in the decision process and that the saccadic burst in SC may be a signature of a threshold computation, which terminates the decision process once SC receives a critical level of input. These results inform our understanding of how networks of brain areas cooperate to form and terminate a decision.

1-079. Hippocampal indexing at the down-to-up transition promotes consolidation without interference

Ryan Golden[1,2] RYGOLDEN@UCSD.EDU
Oscar Gonzalez[1] O2GONZALEZ@UCSD.EDU
Masami Tatsuno[3] TATSUNO@ULETH.CA
Bruce McNaughton[4] BRUCE.MCNAUGHTON@ULETH.CA
Maxim Bazhenov[1] MBAZHENOV@UCSD.EDU

[1]University of California, San Diego
[2]Neurosciences Graduate Program
[3]University of Lethbridge
[4]University of California, Irvine

Systems Consolidation Theory posits that the hippocampus encodes new information for later cortical consolidation during sleep. Specifically, coupling of slow oscillations (SOs) in the cortex and sharp-wave/ripples (SWRs) in the hippocampus is thought to allow the hippocampus to replay recent memories and to index corresponding cortical memory traces to be replayed and learned for long-term storage. To understand the details of this coupling, we analyzed neural ensemble activity from the medial prefrontal cortex (mPFC) and CA1 of the hippocampus from rats trained to run a spatial sequence memory task and developed a biophysically-realistic thalamocortical network model implementing SWR input and SOs. We found in vivo that SWRs resulted in a transient decrease in von Neumann entropy of the cortical spiking activity during UP states, suggesting SWRs can bias activity patterns in mPFC. Despite the phase-independence of this effect, we found a preferred phase for SWRs just after the down-to-up transition (DUt) of SOs. Using our model, we next found that when two competing memories were trained sequentially during awake, the model suffered from retroactive interference, forgetting the old memory trace. However, interference could be avoided when the competing new memory was embedded to the cortical network by SWRs during sleep. This triggered consolidation of the new memory at the same time as reconsolidation of the old one. Importantly, consolidation of the new memory resulted when SWR input occurred at the DUt, but not during the middle of the up state. In this scenario, the new memory was replayed just after the DUt while the old memory was replayed during the middle of the up state. Reminiscent of approaches to overcoming catastrophic forgetting in artificial neural networks, our model predicts that systems consolidation allows the brain to continuously learn by interleaving replay of old and new memories within individual up states.

1-080. Dendritic re-routing accompanying remapping in hippocampal area CA3

Jason Moore[1,2] JASON.MOORE@NYULANGONE.ORG
Jayeeta Basu[1] JAYEETA.BASU@NYULANGONE.ORG

[1]New York University School of Medicine
[2]Neuroscience Institute

Hippocampal “place cells” are tuned to specific locations in an environment. These representations can change or “remap” in different environments or in the same environment in response to changes in task demand. Remapping occurs on fast timescales, so typical plasticity mechanisms cannot completely explain this phenomenon. Dendritic spiking has been implicated in the formation of stable place fields, and here we explore the contributions of dendrites to the flexibility of hippocampal representations. Here we utilize an experimental paradigm in which the basal dendrites, cell bodies, and apical dendrites of pyramidal neurons in hippocampal area CA3 are all imaged in the same focal plane. This allows us to examine a large portion of the dendritic tree, beyond basal or proximal apical dendrites. We imaged calcium activity, reported by GCaMP6f or GCaMP7b, simultaneously in pyramidal cell bodies and dendrites while head-fixed mice ran for water rewards on a treadmill track. To process this data, we developed an automated algorithm utilizing aspects of constrained non-negative matrix factorization (CNMF) to identify dendritic and somatic regions of interest (ROIs) from the calcium signal. The algorithm matches manually-labeled datasets, is easily applied to detecting ROIs from various neuronal morphologies, runs faster than real-time on a standard desktop computer, and has few free parameters to tune. When we analyzed the

compartment-specific tuning of neurons, we discovered that the dendrites of a neuron may be tuned to different locations than their sister dendrites or their parent soma. Furthermore, the dendrites most correlated with the soma often changed during remapping, indicating a change in the internal routing of information. These results suggest that dendrites perform independent computations, and that dendritic outputs are selectively routed to the soma in a context-dependent manner.

1-081. Optimal cue integration in the Drosophila larva

Philip Wong[1] PHILIPWONG@UCSB.EDU
Andreas Braun[2] ANDREAS.BRAUN@PROTONMAIL.COM
Daniel Malagarriga[2] DANIEL.MALAGARRIGA@GMAIL.COM
Jeff Moehlis[1] MOEHLIS@UCSB.EDU
Ruben Moreno Bote[3,4] RUBEN.MORENO@UPF.EDU
Alex Pouget[5] ALEX.POUGET@GMAIL.COM
Matthieu Louis[1] MLOUIS@UCSB.EDU

[1]University of California, Santa Barbara
[2]Center for Genomic Regulation
[3]University Pompeu Fabra, Barcelona
[4]Center for Brain and Cognition
[5]University of Geneva

Numerous studies have established that sensory systems have evolved the capacity to integrate signals (“cues”) from multiple multisensory modalities in a probabilistically optimal fashion. However, these studies have been conducted largely in vertebrates, and while several theories have been proposed to explain how this probabilistic inference can be implemented in large neural populations, it is less clear whether animals with simpler nervous systems like the fruit fly larva are able to implement optimal multisensory integration. To test different mechanistic hypotheses for cue integration, we have developed a behavior-driven agent-based model that reproduces the navigation of Drosophila larvae in sensory gradients.

In this study, we investigated how Drosophila larvae combine sensory cues from real odor, optogenetically-induced virtual odor, and temperature gradients to make navigational decisions in a 2D arena. We constructed an agent-based model that accurately describes the transformation of sensory information into navigational behavior, from running to stopping and turning. By considering a broad set of experimental paradigms, including pairs of sensory modalities with congruent gradients, conflicting gradients, and conditions in which the reliability of sensory information is modulated by optogenetically-induced noise, we find that Drosophila larvae behave in a way consistent with near-optimal cue combination. We report indeed that our mesoscopic agent-based model with Bayesian optimal cue weights provided better fits to the experimental data on average across all conditions compared to two alternative models with winner-take-all strategy (shut weights) or shared fixed weights. In addition, we obtained similar results with a coarse-grained analytical Bayesian optimal model that does not rely on any precise reorientation mechanisms. Altogether, our results suggest that near-optimal cue combination may be present in Drosophila at an early developmental stage. While the neural implementation of probabilistic inference is still unknown, our work suggests that optimal cue combination is an inherent capability of insect nervous systems.

1-082. Activation and disruption of a neural mechanism for novel choice in monkeys

Alessandro Bongioanni[1,2] ALESSANDRO.BONGIOANNI@GMAIL.COM
Miriam Klein-Flugge[3] MIRIAM.KLEIN-FLUGGE@PSY.OX.AC.UK
Matthew Rushworth[3] MATTHEW.RUSHWORTH@PSY.OX.AC.UK

[1]Wellcome Centre for Integrative Neuroimaging, University of Oxford
[2]Experimental Psychology
[3]University of Oxford

The neural circuits underlying the ability to make value-guided decisions have received substantial attention in both humans and animal models, but with conflicting results. Human imaging studies have consistently found effects of subjective value and choice in medial frontal cortex (MFC, sometimes referred to as vmPFC), but electrophysiological experiments in macaques have generally found weak signal there and have focused on orbitofrontal cortex (OFC) as the key area for value representation and decision-making, consistent with lesion studies. A key difference is that animal experiments typically involve long training periods; by contrast, human experiments require choices among options without a history of reinforcement. For example, choosing between two well-known roads both getting to your workplace differs from choosing between two unfamiliar roads when on holidays. We combined imaging in behaving monkeys with neurostimulation and cognitive modelling to uncover a neural system for novel inferential choices. Macaques made binary choices inferring the values of new options via similarities with component parts of familiar options. A whole-brain search with functional magnetic resonance imaging (fMRI) indicated this ability was mediated by MFC, while confirming the role of OFC in overtrained choices. We then investigated the representations allowing this behaviour. OFC encoded the individual identities of the valuable stimuli, but MFC employed a grid-cell-like coding scheme, well known in the context of spatial navigation, to represent the relations of all such stimuli in the abstract space determined by reward amount and probability. Finally, we probed whether this MFC grid-like representation was instrumental in dimension integration during novel choice. We modelled subjective value as a mixture of optimal (multiplicative) and heuristic (additive) strategies for multi-attribute choice. Focused transcranial ultrasound stimulation in MFC, but not adjacent tissue, impaired optimal integration, proving the causal role of the MFC system in novel decision-making.

1-083. Distinct place cell dynamics in CA1 and CA3 encode experience in new contexts

Can Dong[1,2] CANDONG123@UCHICAGO.EDU
Antoine Madar[3] MADAR@UCHICAGO.EDU
Mark Sheffield[3] SHEFFIELD@UCHICAGO.EDU

[1]The University of Chicago
[2]Neurobiology
[3]University of Chicago

A striking feature of the hippocampus is to rapidly generate and continuously update representations of ongoing experience. In particular, representations of contextual experience are known to be encoded by place cells in hippocampal CA1, CA3, and dentate gyrus. Synaptic plasticity within these sub-regions is a key process that alters and stores representations and comes in many forms. However, the synaptic mechanisms leading to place field formation and updating remain disputable. This is in part because the dynamics of the place field emergence and stabilization during the first moments of a novel experience and throughout the ongoing experiences have not been well characterized and systematically compared across different subregions. We therefore used 2-photon calcium imaging in mice running in a novel virtual linear track and tracked the same cells across days to investigate the lap-by-lap emergence dynamics of spatial representations in both CA1 and CA3. We found that place fields appear faster in CA1 but are constantly renewed throughout experience across days whereas they appear later in CA3 but remain relatively stable with experience and across days. After emergence, the location of the place field is not always stable, sometimes showing prominent backward or forward shifting from lap to lap. During the first laps of exploration of a new environment, the average spatial representation in the hippocampus shifts backward over time, with a faster shift in CA1 than CA3 that slows down with familiarization across days. To explain those differences and investigate the mechanisms supporting place field emergence and shifting dynamics in the hippocampus, we plan to use computational models of spiking neurons with plastic synapses following different spike-timing dependency rules. Our goal is to evaluate how differences in connectivity, excitability or long-term plasticity impact lap-by-lap place field dynamics in the hippocampus.

1-084. The virtual patch clamp: imputing C. elegans membrane potentials from calcium imaging

Andrew Warrington[1,2] ANDREW.WARRINGTON@KEBLE.OX.AC.UK
Arthur Spencer[3] ARTHUR.SPENCER@BRISTOL.AC.UK
Frank Wood[4] FWOOD@CS.UBC.CA

[1]University of Oxford
[2]Engineering Science
[3]University of Bristol
[4]University of British Columbia

We tackle whole-connectome scale simulation and inference in a model of the nematode roundworm Caenorhabditis elegans (C. elegant). The main contribution of this work is to present a method for inferring unobserved physiologically relevant states and parameter values from partial data, of the type and dimension that are currently obtainable using calcium fluorescence imaging. This allows previously unobtainable variables of interest, such as membrane potentials, to be estimated non-invasively at whole-connectome scale.

To achieve this, we develop a stochastic C. elegans brain-body simulator, operating at the resolution of individual neurons. We use this simulator in a particle-based posterior inference method to simultaneously estimate the distribution over unobserved variables and parameter values by maximizing the model evidence given the ob-

served data. We demonstrate our approach by successfully recovering latent state trajectories and parameters from synthetic data.

This is the first attempt we know of to "complete the circle," where an anatomically grounded whole-connectome simulator is used to impute a time-varying "brain" state at single-cell fidelity from covariates that are measurable in practice. Our method is highly extensible to incorporate new models and observation modalities as they are developed. We believe this approach is an ambitious first step towards linking low-level neural activity to organism-level behaviors. If applied to whole assays of real organisms, this would offer an unprecedented method for inspecting the neural activity underlying complex behaviors and responses to stimuli. Further, it would provide a method for in silico hypothesis testing using posterior-predictive inference across, potentially, millions of digital specimens, each of which can be perfectly observed at cellular resolution.

1-085. Synthesizing Preferred Stimuli for Individual Voxels in the Human Visual System

Johannes Roth[1] JOHANNES@ROTH24.DE
Katja Seeliger[2] KSEELIGER@POSTEO.JP
Thomas Schmid[1] SCHMID@INFORMATIK.UNI-LEIPZIG.DE
Martin Hebart[3] HEBART@CBS.MPG.DE

[1]University of Leipzig
[2]Max Planck Institute for Human Cognitive and Brain Sciences
[3]Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig

A central aim of visual neuroimaging is to determine the functional selectivity of individual voxels in visual cortex, yet presenting all possible stimuli is impractical given the complexity of the visual world. To overcome this limitation, we developed a convolutional neural network model of occipitotemporal cortex based on a large-scale functional MRI dataset that allowed us to predict voxel responses for novel stimuli in-silico. To identify preferred stimuli for individual voxels, we developed an interpretability technique based on a generative adversarial neural network (GAN) that allowed us to use gradient ascent for synthesizing naturalistic preferred images. As expected, voxels in areas V1-V3 yielded small receptive fields with a preference for gratings. Higher order areas showed mixed preference: While confirming face-selectivity in FFA and place-selectivity in PPA, both regions additionally showed preference for other visual features, such as oval shapes and vertical lines in FFA, or horizontal lines and high spatial frequency in PPA. The GAN latent vectors for FFA and PPA representations were highly distinct (SVM classification accuracy: 87.0%), underscoring the validity of the approach. Together, this approach opens the avenue towards precision functional mapping of selectivity at the level of individual voxels across the whole visual system and thus may reveal previously unknown functional selectivity.

1-086. Beyond the Euclidean brain: inferring non-Euclidean latent trajectories from spike trains

Kristopher Jensen[1,2] KTJ21@CAM.AC.UK
David Liu*[1] DL543@CAM.AC.UK
Ta-Chu Kao[1] TCK29@CAM.AC.UK
Marco Tripodi[3] MTRIPODI@MRC-LMB.CAM.AC.UK
Mate Lengyel[1,2] M.LENGYEL@ENG.CAM.AC.UK
Guillaume Hennequin[1,2] G.HENNEQUIN@ENG.CAM.AC.UK

[1]University of Cambridge
[2]Department of Engineering
[3]MRC Laboratory of Molecular Biology

Neuroscience faces a growing need for scalable data analysis methods that reduce the dimensionality of population recordings yet retain key aspects of the computation or behaviour. To extract interpretable latent trajectories from neural data, it is critical to embrace the inherent topology of the features of interest: head direction evolves on a ring or torus, 3D body rotations on the special orthogonal group, and navigation is best described in the intrinsic coordinates of the environment. Accordingly, the manifold Gaussian process latent variable model (mGPLVM) was recently proposed to simultaneously infer latent representations on non-Euclidean manifolds and how neurons are tuned to these representations. This probabilistic method generalizes previous Euclidean models and allows principled selection between candidate latent topologies. While powerful, mGPLVM makes two unjustified approximations that limit its practical applicability to neural datasets. First, consecutive latent states are assumed independent a priori, whereas behaviour is continuous in time. Second, its Gaussian noise model is inappropriate for positive integer spike counts. Previous work in Euclidean LVMs such as GPFA has shown significant improve-

ments in performance when modeling such features appropriately. Here, we extend mGPLVM by incorporating temporally continuous priors over latent states and flexible count-based noise models. This improves inference on synthetic data, avoiding negative spike count predictions and discontinuous jumps in latent trajectories. On real data, we also mitigate these pathologies while improving model fit compared to the original mGPLVM formulation. In summary, our extended mGPLVM provides a widely applicable tool for inferring (non-)Euclidean neural representations from large-scale, heterogeneous population recordings. We provide an efficient GPU implementation in python, relying on recent advances in approximate inference to e.g. fit 10,000 time bins of recording for 100 neurons in five minutes.

1-087. Attention modulates representation of visual luminance contrast in a layer- and cell-class specific way

Xiang Wang[1,2] XIANG.WANG@YALE.EDU
Anirvan Nandy[3] ANIRVAN.NANDY@YALE.EDU
Monika Jadi[3] MONIKA.JADI@YALE.EDU

[1]Yale University
[2]Interdepartmental Neuroscience Program
[3]Yale School of Medicine

The encoding of contrast information is one of the most important computations performed by visual circuits and is modulated by visual spatial attention signals. It is unknown whether the attentional modulations of contrast computations are homogeneous in the cortex or specific to subpopulations within individual layers. To investigate this question, we recorded the neural responses in visual area V4 to Gabor stimuli with different contrasts while two macaques performed an attention-demanding orientation change detection task. We classified well-isolated single units into five distinct classes by applying an unsupervised machine learning algorithm on two electrophysiological properties extracted from extracellular recordings: peak-to-trough duration and local variation in inter-spike intervals. We found that attentional modulation of contrast responses is cell-class- and layer-specific. Narrow-spiking units (putative interneurons) exhibited contrast-independent attentional modulations (response gain) across all three layers. In contrast, attentional modulations of broad-regular-spiking neurons (putative pyramidal cells) were more robust in the low-contrast range in the superficial and input layers (a mixture of response gain and contrast gain), but were contrast-independent in the deep layer. When interpreted within the framework of the normalization model of attention, the layer-specific contrast dependencies of attentional modulations of broad-regular-spiking units predict broader spatial pooling of local inhibitory populations down the cortical depth. Such differences further predict a layer-specific signature of cross-correlations between the activities of local inhibitory and putative projection neurons when explored in a spiking E-I network model. We found the same patterns of cross-correlations between narrow- and regular-broad-spiking clusters across layers in the experimental data. Our analyses revealed that the attention effects on contrast computations are layer- and cell-class-specific, and the layer-wise differences in attentional modulation patterns may result from variations in inhibitory spatial pooling size across layers.

1-088. Biophysically grounded mean-field models of neural populations under electrical stimulation

Caglar Cakan[1,2] CAKAN@NI.TU-BERLIN.DE
Klaus Obermayer[1] KLAUS.OBERMAYER@TU-BERLIN.DE

[1]Technische Universitat Berlin
[2]Neural Information Processing Group

Electrical stimulation of neural systems is a key tool for understanding neural dynamics and ultimately for developing clinical treatments. Many applications of electrical stimulation affect large populations of neurons. However, computational models of large networks of spiking neurons are inherently hard to study and an understanding of how neural populations interact with electrical inputs is lacking.

We present a reduced mean-field model of excitatory and inhibitory adaptive exponential integrate-and-fire (AdEx) neurons which can be used to efficiently study the effects of electrical stimulation on large neural populations. The rich dynamical properties of this basic cortical model are described in detail and validated using large network simulations. Bifurcation diagrams reflecting the network's state reveal asynchronous up- and down-states, bistable regimes, and oscillatory regions corresponding to fast excitation-inhibition and slow excitation-adaptation feedback loops. The biophysical parameters of the AdEx neuron can be coupled to an electric field with realistic field strengths which then can be propagated up to the population description.

We show how on the edge of bifurcation, direct electrical inputs cause network state transitions, such as turning on and off oscillations of the population rate. Oscillatory input can frequency-entrain and phase-lock endogenous oscillations. Relatively weak electric field strengths on the order of 1 V/m are able to produce these effects, indicating that field effects are strongly amplified in the network. The effects of time-varying external stimulation are well-predicted by the mean-field model.

Our results show that 1) weak fields of 1 V/m are able to affect neural population activity, as was observed experimentally 2) the response to stimulation critically depends on the state of the neural system and 3) mean-field models offer an efficient and appropriate framework for investigating the effects of electric fields on a cortical system.

1-089. Robust representation of Natural Scenes within Columns of Primary Visual Cortex

Xiaomo Chen[1,2] XIAOMO@STANFORD.EDU
Kunlun Bai[1] BAIKUNLUN1023@GMAIL.COM
Shude Zhu[1] SHUDE@STANFORD.EDU
Ruobing Xia[1] RUOBINGX@STANFORD.EDU
Nathan Kong[1] NCLKONG@STANFORD.EDU
Sebastian Weingartner[1] S.WEINGARTNER@TUDELFT.NL
Anthony Norcia[1] AMNORCIA@STANFORD.EDU
Tirin Moore[1] TIRIN@STANFORD.EDU

[1]Stanford University
[2]Neurobiology

The visual system is believed to have adapted to the statistical properties of the natural environment. However, the stage at which sensitivity to naturalistic structure emerges during visual processing is controversial. To address this, we recorded in macaque V1 using high-density electrode arrays (Neuropixels probes, IMEC) and investigated the neural representation of images in macaque primary visual cortex (V1) using three sets of images that varied in their higher order image statistics: natural scenes (NA), synthetic natural texture (ST) images, and phase-scrambled synthetic noise (SN) images. We measured the response from hundreds of V1 neurons across the columnar structure of V1. We first examined the sensitivity of single neurons for natural images and found that V1 neurons responded more vigorously to NA than to SN and ST images within the first 50 ms of the visual response. This sensitivity to natural images structure was robustly observed throughout all cortical layers, including neurons in the input layer 4C. We next examined the population coding of V1 neurons for natural scenes. We found that V1 neurons showed high population sparseness and high lifetime sparseness for natural scenes. Moreover, individual neurons were less synchronized with the overall population activity for images containing more natural statistical components. In addition, the neural ensembles showed higher discriminability for natural scenes. Lastly, we examined the activity patterns of an ImageNet-trained deep convolutional neural network model (DCNN) in response to the same set of stimuli. This analysis shows that higher sensitivity to natural scenes also emerges early in this feedforward artificial visual system. Taken together, these results revealed a distinct coding preference for natural images starting at the earliest stages of visual cortex. In addition to feedback modulation from higher visual cortex, these results suggest a feed-forward sensitivity to natural scenes during cortical visual processing.

1-090. Distinct neural mechanisms construct classical versus extraclassical inhibitory surrounds in an inhibitory nucleus in the midbrain attention network

Hannah Schryver[1,2] HANNAH.SCHRYVER@JHU.EDU
Shreesh Mysore[1] MYSORE@JHU.EDU
Jing-Xuan Lim[3] JINGXLIM@JHMI.EDU

[1]Johns Hopkins University
[2]Psychological and Brain Sciences
[3]HHMI Janelia Research Campus

Selecting the location with the most important stimulus to guide behavior at any instant is an essential part of adaptive behavior. The superior colliculus (SC, or optic tectum OT, in mammals) is required for such spatial target selection, with competitive interactions among its representations of stimuli being critical [1, 2]. In turn, in birds, inhibitory neurons external to the OT, called Imc (isthmi pars magnocellularis) [3-5] control these competitive interactions and stimulus selection by the OT [4, 6-8]. A core characteristic of neurons involved in spatial selection that impacts both their ability to represent preferred stimuli selectively and compare competing stimuli is their

spatial receptive field (RF). The excitatory center and classical inhibitory surround of the RF together control the responses of neurons to a stimulus inside the RF, whereas the extraclassical surround controls the modulation of the neuron's responses by competing stimuli outside. Here, we discover in the barn owl how classical as well as extraclassical (global) inhibitory surrounds of Imc receptive fields (RFs), fundamental units of Imc computational function, are constructed. We first found that focal, reversible (iontophoretic) blockade of GABAergic input onto Imc neurons disconnects their extraclassical inhibitory surrounds but, surprisingly, leaves intact their classical surrounds. To unpack this intriguing finding, we made paired recordings at spatially aligned sites in the Imc and OT. We found that iontophoretically blocking inhibition onto OT10 neurons – the sole source of excitatory input to Imc neurons – disrupts classical inhibitory surrounds of Imc RFs. Separately, with paired recordings at sites within Imc encoding for distant locations, we found that iontophoretically silencing distant neurons abolished global competitive surrounds of Imc RFs. Thus, whereas classical inhibitory surrounds of Imc neurons are inherited from upstream OT, their extraclassical inhibitory surrounds are constructed locally within Imc. These results reveal fundamental design principles of the midbrain spatial attention circuit and attest to the critical importance of competitive interactions within Imc for its operation.

1-091. Distinct synaptic plasticity timescales underlie engram dynamics in systems consolidation

Douglas Feitosa Tome[1,2] DOUGLAS.FEITOSA-TOME17@IMPERIAL.AC.UK
Sadra Sadeh[1] S.SADEH14@IMPERIAL.AC.UK
Claudia Clopath[1] C.CLOPATH@IMPERIAL.AC.UK

[1]Imperial College London
[2]Bioengineering

Previous experiments have shown that the initial state and subsequent evolution of engram cells in systems consolidation are specific to each brain region. It has been proposed that the observed engram dynamics are characteristic of generic episodic memories. However, the exact mechanisms underlying engram cell dynamics have not been investigated yet. Here, we present a spiking network model of systems consolidation that captures the behavior of engram cells in cortex and hippocampus. Our model contains individual subnetworks for the hippocampus, thalamus, and cortex with excitatory and inhibitory synapses within and between these regions being subject to plasticity mechanisms at different rates. In this model, cortical engram cells are initially silent in recent recall but mature and become active in remote recall after a consolidation period. Inversely, hippocampal engrams are active in recent but not in remote recall and, hence, de-mature. Hippocampal engram cells as well as spontaneous coordinated reactivations in the three brain regions in our model were crucial for systems consolidation of memories. Importantly, the engram dynamics and interdependencies displayed by our model are aligned with previous experimental findings. Our model also makes the following experimentally testable predictions: engram cells in mediodorsal thalamus are active in recent and remote recall and are essential for the consolidation of prefrontal cortex engram cells, inhibitory engram cells have distinct dynamics with coupled reactivations, and the synaptic coupling between mediodorsal thalamus and prefrontal cortex prior to consolidation is predictive of the retrograde amnesia curve induced by hippocampal damage. Overall, our modelling results suggest that a variety of synaptic plasticity mechanisms operating over distinct timescales across brain regions is coordinated to promote systems consolidation.

1-092. The arm's posture does not alter the time course of population activity in motor cortex

Patrick Marino[1,2] PMARINO162@GMAIL.COM
Emily Oby[1] EMILYOBY@GMAIL.COM
Asma Motiwala[3] AMOTIWALA@CMU.EDU
Erinn Grigsby[1] ERINN.GRIGSBY@GMAIL.COM
Alan Degenhart[4] ALANDEGENHART@GMAIL.COM
Byron Yu[3] BYRONYU@CMU.EDU
Aaron Batista[1] AARON.BATISTA@GMAIL.COM

[1]University of Pittsburgh
[2]Bioengineering
[3]Carnegie Mellon University
[4]Allen Institute for Brain Science

Dynamical structure has been observed in neural activity in the primary motor cortex (M1) during reaching movements. This indicates that M1 might function as a dynamical system in which patterns of activity are generated

through the interaction of local dynamics and external inputs. An important open question is how external inputs shape the time course of population activity in M1.

Postural signals are a major source of input to M1. It is known from classic single-neuron studies that altering arm geometry during reaching movements strongly affects neural tuning and firing rate range. However, previous experiments do not isolate the effect of postural inputs on M1's dynamics, because changing arm geometry during movement also changes muscle activity. Thus, changes in neural activity observed during movement cannot be attributed to postural inputs alone.

Here, we leveraged a brain-computer interface (BCI) to investigate the effects of postural inputs on M1's dynamics. Using a BCI, we can change arm posture without changing muscle activity during the neural control of an effector. To change postural inputs to M1, we systematically rotated a monkey's arm about the shoulder while the animal controlled the BCI. We found that M1 encoded the posture of the arm in a low-dimensional subspace which was orthogonal to the subspace typically occupied by neural activity during BCI control. The posture of the arm could be decoded from this subspace regardless of BCI target direction. Neural trajectories evolved from posture-specific initial conditions. Surprisingly, these changes in initial condition did not change the temporal evolution of the trajectories. These properties may facilitate two capacities of the motor system: first, circuits involved in movement control can access information about the posture of the body, regardless of the behavior in which the body is engaged. Second, postural information can be used flexibly in guiding movements.

1-093. Cortical mechanisms of saccadic suppression

Sachira Denagamage[1] SACHIRA.DENAGAMAGE@YALE.EDU
Mitchell Morton[1,2] MITCHELL.MORTON@YALE.EDU
John Reynolds[3] REYNOLDS@SALK.EDU
Monika Jadi[1] MONIKA.JADI@YALE.EDU
Anirvan Nandy[1] ANIRVAN.NANDY@YALE.EDU

[1]Yale School of Medicine
[2]Neuroscience
[3]Salk Institute for Biological Studies

Saccades are a ubiquitous and crucial component of our visual system, allowing for the efficient deployment of the fovea and its accompanying neural resources. Initiation of a saccade is known to cause saccadic suppression, a temporary reduction in visual sensitivity and visual cortical firing rates. While saccadic suppression has been well characterized at the level of perception and single neurons, relatively little is known about the visual cortical networks governing this phenomenon. Here we examine the effects of saccadic suppression on distinct neural subpopulations within area V4 of the macaque. We find cortical layer- and cell type-specific differences in the magnitude and timing of peri-saccadic modulation. Neurons in the input layer show changes in firing rate and inter-neuronal correlations prior to saccade onset, suggesting that this layer receives information about impending saccades. Putative inhibitory interneurons in the input layer elevate their firing rate during saccades and may play a role in suppressing the activity of other cortical subpopulations. A computational model of this circuit recapitulates our empirical observations and demonstrates that an input layer-targeting pathway could initiate saccadic suppression by enhancing local inhibitory activity. Together, our results provide a mechanistic understanding of how eye movement signals interact with cortical circuitry to enforce saccadic suppression.

1-094. Transformation of population representations of sounds throughout the auditory system

Jacques Bourg[1] JBOURG@PASTEUR.FR
Sophie Bagur[2] SOPHIE.BAGUR@PASTEUR.FR
Alexandre Kempf[3] KEMPF@INAF.CNRS-GIF.FR
Thibault Tarpin[3] THIBAULT.TARPIN@ENS.FR
Khalil Bergaoui[3] KHALIL.BERGAOUI@STUDENT.ECP.FR
Yin Guo[4] YIN.GUO@PASTEUR.FR
Etienne Gosselin[4] ETIENNE.GOSSELIN@PASTEUR.FR
Jean-Luc Puel[5] JEAN-LUC.PUEL@INSERM.FR
Jerome Bourien[5] JEROME.BOURIEN@UMONTPELLIER.FR
Sebastian Ceballo Charpentier[2] SEBASTIAN.CEBALLO-CHARPENTIER@PASTEUR.FR
Brice Bathellier[2] BRICE.BATHELLIER@PASTEUR.FR

[1]Hearing Institute, Institut Pasteur 63 rue de Charenton
[2]Hearing Institute, Institut Pasteur and CNRS
[3]CNRS

[4]Hearing Institute, Institut Pasteur
[5]Institute for Neurosciences of Montpellier, INSERM - UMR 1051, University of Montpellier, Montpellie

Auditory perception relies on robust spectrotemporal pattern separation. Yet the step-by-step computational transformations through which this occurs in the auditory system are still elusive, partly due to the lack of systematic comparison of sound representations at different stages. To address this, we compared population representations based on extensive two-photon calcium imaging and electrophysiology recordings in the auditory cortex, thalamus, inferior colliculus in mice and on a detailed biophysical model of the cochlea. Using a noise-corrected correlation metric, we measured the similarity of activity patterns generated by diverse canonical spectrotemporal motifs (pure tones and chords, amplitude and frequency modulations, broad band noise). This measure evidenced a gradual decorrelation of population responses across sounds, accompanied by a convergence of the information carried by temporal and rate codes in cortex. While this transformation was related to response sparsenning and strong non-linearities, the sound representations remained continuous, preserving local similarity relationships at all stages. Decorrelation rate depended on input similarity. Rate code representations that were most similar at cochlear level, for example for time symmetric sounds were decorrelated mainly between thalamus and cortex, while many other sound pairs underwent significant decorrelation earlier. Strikingly, representations of different intensities of the same sounds were also decorrelated. This feature and gradual decorrelation of time-independent representations could be reproduced by a deep network trained to detect in parallel basic perceptual attributes of sounds, including estimates of frequency and intensity ranges as well as temporal modulation types. In contrast, networks trained on tasks ignoring some perceptual attributes (e.g. intensity) did not decorrelate the related acoustic information. Together these results establish that the mouse auditory system transforms sound representations to make information about diverse perceptual attributes generically accessible in a rate-based population code.

1-095. A continuous information theoretic method to estimate millisecond scale timing precision

Tobias Niebur[1,2] TNIEBUR3@GATECH.EDU
Joy Putney[1] JPUTNEY3@GATECH.EDU
Simon Sponberg[1] SPONBERG@GATECH.EDU

[1]Georgia Institute of Technology
[2]Electrical and Computer Engineering

Precise timing codes are prevalent in sensory systems, but we are just beginning to appreciate their commonality in motor systems. In order to assess how a motor program is encoded and coordinated in timing, it is necessary to characterize the precision of each motor unit. A prior method for assessing precision utilized binning of spike times and direct estimation of mutual information. However, this fails at very high precision, where low sampling of spike train states at small bin sizes causes instability. To rectify this, we extend the Kraskov continuous information theoretic estimation technique by progressively corrupting the spike times with additive uniform noise. We apply this to the recently recorded, comprehensive, spike-resolved motor program of hawkmoth flight muscles and demonstrate millisecond-scale precision exists in every motor unit.

1-096. Differential signatures of top-down and bottom-up biases in auditory decision-making

Nathan Tardiff NTARDIFF@SAS.UPENN.EDU
Lalitta Suriya-Arunroj LALITTA.SURIYA.ARUNROJ@GMAIL.COM
Yale University Cohen YCOHEN@PENNMEDICINE.UPENN.EDU
Joshua Gold JIGOLD@PENNMEDICINE.UPENN.EDU

University of Pennsylvania

Perceptual decisions can be affected by different sources of prior information. For example, top-down cues that explicitly indicate stimulus probabilities tend to bias decisions toward more-probable stimuli. Alternatively, bottom-up cues in the form of stimulus regularities can cause habituation to repeated stimuli, which sensitizes perceptual systems to violations of those regularities. However, how these top-down and bottom-up effects interact to affect perceptual decisions, particularly for auditory stimuli that can be affected strongly by habituation, is not well understood. To address this knowledge gap, we compared how top-down and bottom-up cues, presented separately or together, biased the choices and response times (RTs) of 50 human subjects performing a frequency-discrimination task. Subjects were presented with a temporal sequence of tone bursts (200 or 2500 Hz; 300 ms tone duration) and then reported the frequency (“low” or “high”) of the final “test” tone burst. Task difficulty was

titrated by embedding the test tone in a background of white noise using varying signal-to-noise ratios (SNRs). Top-down cues induced choice and RT biases toward the more-probable stimulus, replicating previous results. Bottom-up cues had a similar influence when the SNR of the test stimulus was low, acting like a prior that biased choice and RT toward the frequency of the tones presented just before the test stimulus. In contrast, bottom-up cues had a different, adaptation-like effect when the SNR of the test stimulus was high, biasing choice and RT away from the frequency of the recently presented tones. We captured these effects using a novel drift-diffusion model (DDM) with two types of bias: 1) an additive, expectation-congruent effect originating from either top-down or bottom-up cues; and 2) an SNR-dependent adaptation effect on perceptual sensitivity, modulated by bottom-up signals. Our results provide a new, unified perspective on how top-down and bottom-up information jointly affect auditory decisions.

1-097. A timeline of past sensory and future motor events in the anterior lateral motor cortex

Rifqi Affan[1,2] RIFQIAFFAN@GMAIL.COM
Ian Bright[1] IMBRIGHT@BU.EDU
Marc Howard[1] MARC777@BU.EDU

[1]Boston University
[2]Psychological and Brain Science

Temporal information about past events can be represented in the brain through the sequential activations of “time cells” in the hippocampus or the monotonic decay of neurons in the entorhinal cortex (EC)1. Monotonic activity of neurons observed in the mouse anterior lateral motor (ALM) cortex resembles the dynamics found in EC but is instead tuned to the direction of planned future movements2. However, it is unclear whether and how neurons in the ALM cortex carry information about time until a planned movement. Here, we reanalyze a publicly available dataset provided by Inagaki and colleagues2 on https://crcns.org/data-sets/motor-cortex/alm-5. The data contains electrophysiological recordings of ALM neurons in mice during an auditory delayed-response task. We use a model-based approach to quantify the temporal response of each neuron from 0.5 s prior to the onset of auditory cue presentation to 1 s after a “Go” cue. Single-neuron responses are fit with exponentially modified Gaussian functions, providing estimations of peak time, peak amplitude, and time constant of ramp/decay. Likelihood ratio analysis reveals two classes of time-varying neurons better fit by the ex Gaussian model compared to the gaussian model, (df = 1, p < 0.01 Bonferroni corrected). One class of ALM neurons peak at the start of cue period then decays to baseline (36 of 562 time-varying, putative pyramidal neurons). The other class of neurons ramp up to a peak around the “Go” cue (126/562). Both the ramping and decay dynamics across ALM neurons exhibit similar peak times but diverse time constants, unlike sequentially-activated cells that show different peak times. Cross-temporal decoder analysis shows that the direction of future movement can be decoded from ALM population activity, with an accuracy gradient across time. Our results indicate that monotonic dynamics in the ALM cortex carry temporal information through a wide range of time constants.

1-098. Division of labor among interneurons enables rich cortical computations

Gregory Handy[1,2] GHANDY@UCHICAGO.EDU
Destinee Aponte[3] DESTINEE_APONTE@G.HARVARD.EDU
Hiroyuki Kato[4] HIROYUKI_KATO@MED.UNC.EDU
Brent Doiron[1] BDOIRON@UCHICAGO.EDU

[1]University of Chicago
[2]Neurobiology
[3]Harvard University
[4]University of North Carolina

Interneurons play a key role in many components of sensory processing, including modulating feature selectivity and maintaining an asynchronous network state. Past mathematical models have furthered our understanding of some of these processes by considering a recurrent network of excitatory and inhibitory neurons. However, experimental evidence has shown that diversity exists within this inhibitory population, with 70% of interneurons falling into one of two major subtypes: parvalbumin-expressing (PV) and somatostatin-expressing (SOM) neurons. New models are required to parse how diverse inhibitory circuits contribute to cortical processing.

The differences in how inhibitory subtypes are embedded into the cortical circuit are considerable. For example, in the primary auditory cortex (A1) of mice, SOM neurons project and receive spatially broader connections than other inhibitory subtypes. While this suggests that inhibitory subtypes are primed to play particular roles in driving

network dynamics, these exact roles are unknown. Here, we investigate how individualized roles within the circuit enables neurons to selectively respond to directional auditory sweeps. Using a combination of imaging and whole-cell recording techniques we show that low frequency preferring neurons selectivity respond to upward frequency sweeps, while high frequency preferring neurons respond to downward sweeps. Optogenetically inhibiting PV interneurons yielded minimal changes to direction selectivity, while inhibiting SOM interneurons significantly reduced it.

To further investigate the circuit mechanisms responsible for this finding, we construct a mathematical model of A1 which includes spatially distributed coupling over the tonotopy. Similar to our experiments, this model finds that SOM neurons are the key subpopulation in establishing direction selectivity. The model is then used to demonstrate that this ability arises due to their broad spatial projections. We also prove that a nonlinear transfer function is a necessary component of the model, whereas a simplified linearized model would fail to exhibit direction selectivity.

1-099. Urgency as the neural correlate of the opportunity cost of time commitment in decision making tasks

Maximilian Puelma Touzel[1] MAX.PUELMATOUZEL@GMAIL.COM
Paul Cisek[2] PAUL.CISEK@UMONTREAL.CA
Guillaume Lajoie[3] G.LAJOIE@UMONTREAL.CA

[1]Universite de Montreal, Montreal Institute for Learning Algorithms
[2]Universite de Montreal
[3]Universite de Montreal, Mila

Humans and animals can make decisions that incorporate the relative value of alternatives. For example, a student taking an exam may weigh the benefit of spending more time on a difficult question versus moving on to easier ones. An inevitable cost of moving on is the value of the unrealized alternative, i.e. of eventually succeeding on the hard question. This value forfeited by selecting a given action is known as opportunity cost. Importantly, decision-making tasks used in neuroscience typically involve opportunity costs. However, their role in shaping decision-making and their neural activity correlates are unknown. In parallel, urgency, a ramping signal in trial-aligned neural activity, is thought to control the speed of decision-making. What determines its shape and why it differs qualitatively across tasks is unknown. Here, we propose that urgency is the neural correlate of opportunity cost and provide a theory, validated with published primate recordings, that interpretably dictates its shape depending on task demands. Specifically, we show (1) a linearly rising component associated with time-averaged reward rate, (2) a context-dependent offset, and (3) a within-trial saturation arising from inferring expected trial reward.

These results use a modification of average reward reinforcement learning, previously used to propose the static time-averaged reward as the rate at which opportunity cost is incurred. Realistically however, the opportunity cost rate is dynamic on multiple timescales. We extend AR-RL to this case by recasting the average as a time-varying belief average over the combined opportunity costs of time commitment. We show that adaptive or multi-channel filtering of reward prediction errors form a reward representation with which to construct the necessary modifications of the AR-RL problem. By summarizing complex environment dynamics, urgency supports learning that straddles model-based and model-free methods and motivates a multiscale effort to further dissect the neural basis of relative-value decision making.

1-100. Evolving to learn: discovering interpretable plasticity rules for spiking networks

Jakob Jordan[1,2] JAKOB.JORDAN@UNIBE.CH
Maximilian Schmidt[3] SCHMIDT.MAXIMILIAN@POSTEO.DE
Walter Senn[4] WALTER.SENN@UNIBE.CH
Mihai Petrovici[5,6] MIHAI.PETROVICI@UNIBE.CH

[1]University of Bern
[2]Department of Physiology
[3]RIKEN Center for Brain Science
[4]Universitat Bern
[5]University of Bern and Heidelberg University
[6]Department of Physiology and Kirchhoff-Institute for Physics

How do we learn? Whether we are memorizing the way to a lecture hall or mastering a sport, our central nervous system is able to retain the relevant information over extended periods of time. Adjustments of the interaction strength between neurons are a key component in this process. How these can be mathematically described at the phenomenological level, as so called “plasticity rules”, is essential both for understanding biological information processing and for developing cognitively performant artificial systems.

We suggest an automated approach for discovering biophysically plausible plasticity rules. By evolving compact symbolic expressions we ensure the discovered plasticity rules are amenable to intuitive understanding. This is fundamental for successful communication and human-guided generalization, for example to different network architectures or task domains.

We apply out method to three learning paradigms. In a reward-driven learning scenario we demonstrate that, in contrast to the prevailing view (e.g., Fremaux et al., 2010), agents with no estimate of the expected reward outperform agents who make use of such information. In an error-driven learning scenario a set of discovered plasticity rules can be interpreted as variations on the dendritic prediction of somatic spiking (Urbanczik & Senn, 2014). In a correlation-driven learning scenario, the evolutionary search discovers a variety of STDP kernels and associated homeostatic mechanisms that are functionally indistinguishable. This suggests to reconsider the STDP variations reported in the experimental literature from a point of functional equivalence.

We view the presented methods as a machinery for generating, testing, and extending hypotheses on learning in spiking networks driven by problem instances and prior knowledge and constrained by experimental evidence. We believe this approach can accelerate progress towards deep insights into information processing in physical systems, both biological and biologically inspired, with immanent potential for the development of powerful artificial learning machines.

1-101. Modeling individual variability in neural tuning during auditory perceptual learning

Kathleen Martin[1,2] KAM953@NYU.EDU
Colin Bredenberg[1,3] CJB617@NYU.EDU
Eero Simoncelli[1] EERO.SIMONCELLI@NYU.EDU
Cristina Savin[1,4] CS5360@NYU.EDU
Robert Froemke[1] ROBERT.FROEMKE@NYULANGONE.ORG

[1]New York University
[2]Center for Neural Science
[3]Neural Science
[4]CNS, CDS

Auditory perceptual learning reliably enhances cortical representations of task-relevant stimuli in trained animals relative to naive ones. Despite being associated with perceptual improvement, such changes in neural tuning are typically not measured throughout learning. Additionally, there is a large degree of variability in perceptual learning rates across animals, which is largely ignored when recording and interpreting neural activity. To address these limitations, we developed an experimental and computational framework for describing how sensory representations change during auditory perceptual learning. We recorded from a population of neurons throughout the duration of auditory conditioning using two-photon imaging of layer 2/3 excitatory neurons in the auditory cortex of mice. The animals were progressively trained to classify tones as a single, center frequency or non-center through multiple phases of learning. Perceptual discrimination between center and non-center frequencies improved at variable rates, and across individuals, we observed substantial variability in tuning despite similar behavioral performance. Specifically, animals exhibited either a relative enhancement or suppression of the ‘center’ frequency. To make sense of these observations, we trained a network model by reward-modulated Hebbian synaptic learning (Williams, 1992) to solve the same task. We found that the simulated network learns at similar rates as real animals and captures the across-animal variability in tuning. Overall, both data and model reveal nontrivial learning dynamics associated with perceptual learning in auditory cortex, with initial tuning driving across-animal variability in the emerging representations.

1-102. Outcome dependent modulation of sensory discrimination accuracy by spontaneous cortical state fluctuations

Davide Reato[1] DAVIDE.REATO@NEURO.FCHAMPALIMAUD.ORG
Raphael Steinfeld[2] RAPHAEL.STEINFELD@NEURO.FCHAMPALIMAUD.ORG
Andre Tacao-Monteiro[2] ANDRE.MONTEIRO@NEURO.FCHAMPALIMAUD.ORG
Alfonso Renart[1] ALFONSO.RENART@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Centre for the Unknown
[2]Champalimaud Research, Champalimaud Centre for the Unknown

Desynchronized spontaneous activity (DSA) permits more discriminable sensory representations, but studies in behaving animals suggest that DSA might be linked to engagement or responsivity rather than discrimination accuracy. We studied the effect of two global (face movement and pupil size) and two local (baseline population firing rate (FR) and synchrony (Synch)) state measures on accuracy and engagement in head fixed mice during a frequency discrimination task. Our main result is that baseline fluctuations have a strong impact on discrimination accuracy (roughly doubling the slope of the average psychometric function), but only if the previous trial was an error. Mice were more accurate after errors, but trial outcome did not significantly change baseline activity in the next trial. Furthermore, the outcome-dependent effect of state on accuracy is the same when matching the statistics of baseline FR and Synch after corrects and after errors through trial sub-sampling, suggesting that trial outcome modulates the impact of existing DSA on accuracy, as opposed to affecting DSA and accuracy independently. Movement and baseline firing rate had an effect on engagement (probability of mice not responding to the stimulus), suggesting that mice move while they are distracted from the task. The fraction of variance that FR+Synch explained as to whether a mouse would skip a response was roughly 6 times smaller than the fraction of variance they explained about whether a particular response would be correct after errors, suggesting a stronger impact of baseline fluctuations on accuracy than engagement for our mice. Our results are consistent with DSA in mouse sensory cortex being important for perceptual accuracy provided those cortical areas are made relevant for the task by signals generated by an error, and highlight the importance of accounting for trial history when assessing the behavioral impact of cortical state fluctuations.

1-103. Midbrain dopamine activity during reinforcement learning reflects bias-variance tradeoff

Minryung Song[1,2] RYUNG20@GMAIL.COM
Sue Hee Huh[3] SUEHEEHUH@GMAIL.COM
Jung Whan Shin[4] NEOSJH2009@GMAIL.COM
Min Whan Jung[1] MWJUNG@KAIST.AC.KR
Sang Wan Lee[1] SANGWAN@KAIST.AC.KR

[1]Korea Advanced Institute of Science and Technology
[2]Dept. of bio and brain engineering
[3]Solarluce
[4]Seoul National University Hospital

To maximize reward in a context-changing environment, a reinforcement learning (RL) agent should quickly recognize the changes before adaptation (Fig 1A), which can be facilitated by bias error signaling. Previous studies have shown that, when an animal knows that the environment switches from one familiar setting to another, dopamine reward prediction error (RPE) coding instantly adapts to the change [1]. However, environmental changes are often hidden and have to be inferred. Little is known about dopamine activity while an animal tries to figure out the current context corresponds to which environmental setting it previously learned. We hypothesized that dopamine signals variance (or RPE) in a stable environment, but encodes bias error while the animal in a changing environment is trying to figure out which bias fits the current environmental setting. To test this hypothesis, we analyzed dopamine spike trains recorded while rats were performing a probabilistic two-arm bandit reversal learning task. Previous studies have suggested that animals extensively trained in a reversal task learns the task structure, and use the knowledge to bias animals' choices [2]. We found that extensively trained rats tend to choose a rewarded option consistently but occasionally tested the other option. Dopamine activity during the brief testing resembled bias error. These findings suggest that dopamine activity not only encodes RPE but also signals the bias error in a changing environment. In the broader context, the results also imply that dopamine has the capacity to signal information necessary for adaptation in a volatile environment.

1-104. The spatial structure of feedforward information in mouse primary visual cortex

Jun Zhuang[1,2] ZHUANGJUN1981@GMAIL.COM
Naveen Ouellette[1] NAVEENO@ALLENINSTITUTE.ORG
Rylan Larsen[1] RYLANSL@GMAIL.COM
Kevin Takasaki[1] KEVINT@ALLENINSTITUTE.ORG
Tanya Daigle[1] TANYAD@ALLENINSTITUTE.ORG
Bosiljka Tasic[1] BOSILJKAT@ALLENINSTITUTE.ORG
Hongkui Zheng[1] HONGKUIZ@ALLENINSTITUTE.ORG
Jack Waters[1] JACKW@ALLENINSTITUTE.ORG
R. Clay Reid[1] CLAYR@ALLENINSTITUTE.ORG

[1]Allen Institute for Brain Science
[2]Neural Coding

Location-sensitive and motion-sensitive units are two major functional types of feedforward projections from lateral geniculate nucleus (LGN) to primary visual cortex (V1) in mouse. The distribution of these inputs in cortical depth remains under debate. We systematically mapped the spatial organizations of location-sensitive and motion-sensitive LGN inputs in V1 of awake mice by imaging their axonal calcium activities. Although both types distributed similarly across cortical depth, we found notable organizations specific to each type: the motion-sensitive axons showed a depth-dependent bias in motion direction while the location-sensitive axons showed a depth-independent OFF dominance. By grouping boutons into axon segments, we showed the motion-sensitive axons had greater horizontal bouton spread than location-sensitive axons. Comparisons of response properties between V1 cells and LGN boutons showed that while the retinotopy of V1 cells resembled that from LGN boutons, motion direction preferences of V1 cells were different from their LGN inputs. Overall, our results suggest a new model of parallel thalamocortical pathways: the location information (retinotopy) is relayed from LGN to V1 but the motion information (direction selectivity) may undergo complex computations. These differences are likely mediated by different axon morphologies of these two pathways.

1-105. A Biologically Plausible Neural Network for Circular Kalman Filtering

Anna Kutschireiter[1,2] ANNA.KUTSCHIREITER@GMAIL.COM
Melanie Basnak[1] MELANIEBASNAK@G.HARVARD.EDU
Rachel Wilson[1] RACHEL_WILSON@HMS.HARVARD.EDU
Jan Drugowitsch[1,3] JAN_DRUGOWITSCH@HMS.HARVARD.EDU

[1]Harvard Medical School
[2]Department of Neurobiology
[3]Neurobiology

When navigating through the world, animals need to keep track of dynamic estimates of their position, velocity and heading direction. Some animals' behavior suggests that they also track uncertainties of the variables' estimates, and make strategic use of them. Bayesian inference provides a principled approach to handling unreliable sensory information, yet little is known about how Bayesian algorithms are implemented in the brain's neural networks. Here, we propose a recurrent neural network model that mimics the Drosophila heading direction (HD) system and that tracks both a dynamical HD estimate as well as the uncertainty of this estimate. To develop the network model, we first formulated the sense of direction as a dynamic Bayesian inference task, and derived an approximate closed-form HD tracking algorithm. This "circular Kalman filter" resembles a classic Kalman filter but, due to the circular nature of HD, must invert a nonlinear generative model. Second, we identified a recurrent neural network model that can exactly implement the circular Kalman filter. Importantly, our network architecture matches that of the recently published Drosophila heading direction network. Our model encodes the HD estimate as the peak of a bump of neural activity, and the estimate's certainty as the overall population gain, thereby generalizing standard ring-attractor models that are commonly used to describe biological HD systems. Specifically, we found that the resulting dynamics of the neural population activity are similar to attractor dynamics in the high-certainty regime, but deviates from attractor dynamics when the estimate becomes less certain. Overall, our work demonstrates that the Drosophila HD system could implement a dynamic Bayesian inference algorithm in a biologically plausible manner, and makes specific predictions how this can be tested in experiments.

1-106. A neural dip reveals a motor suppression strategy for changing evidence in perceptual decision-making

Maxwell Shinn[1,2] MAXWELL.SHINN@YALE.EDU
John Murray[1] JOHN.MURRAY@YALE.EDU
Daeyeol Lee[3] DAEYEOL@JHU.EDU
Hyojung Seo*[1] HYOJUNG.SEO@YALE.EDU

[1]Yale University
[2]Psychiatry
[3]Johns Hopkins University

Perceptual decisions based on noisy evidence are thought to utilize a strategy based on the temporal integration of sequentially sampled evidence, which can be formalized through models like the drift-diffusion model. This strategy has been supported by many behavioral studies and is qualitatively consistent with neural activity in multiple brain areas. However, it is unknown what strategies may be used when the strength of sensory evidence changes over time. Here, we trained two monkeys to identify the dominant color of a dynamically refreshed bicolor stimulus and recorded from single neurons in the frontal eye field (FEF). In the task, the stimulus abruptly changed from uninformative to informative during each trial after a variable delay. Immediately after the change in evidence, animals briefly suppressed behavioral responses, creating a dip in the reaction time distribution. Likewise, many neurons in FEF displayed a corresponding dip in activity at this time. This dip has similar form and latency to a dip in activity frequently observed at stimulus onset. We tested three potential strategies by building a generalized drift diffusion model for each strategy and fitting the models to the reaction time distribution. We found that behavioral data, FEF recordings, and microsaccade rate are all well-explained by a brief suppression of motor output, but not by pausing or resetting evidence accumulation. These results show that changes in evidence are associated with distinct patterns of behavior and neural activity, and that these patterns are consistent with motor suppression at the time of onset. This demonstrates how models can provide dissociable predictions to uncover behavioral strategies.

1-107. Network-level computational advantages of single-neuron adaptation

Victor Geadah[1,2] VG351@CAM.AC.UK
Giancarlo Kerg[3] GIANCARLO.KERG@GMAIL.COM
Stefan Horoi[3] STEFAN.HOROI@UMONTREAL.CA
Guy Wolf[3] GUY.WOLF@UMONTREAL.CA
Guillaume Lajoie[4] G.LAJOIE@UMONTREAL.CA

[1]University of Cambridge
[2]Department of Applied Mathematics and Theoretical Physics
[3]Mila, Universite de Montreal
[4]Universite de Montreal, Mila

The impact of single-neuron properties on network-level learning and computation in the brain is not well understood. One such property is the nonlinear transformation of inputs single neurons perform, known as activation functions or nonlinearities. Recent findings highlight how critical the shape of this function is in both dynamical behavior of neural networks and how populations of neurons in the brain encode and process information. Above all, a large literature emphasizes that neuronal activation properties adapt to changes in stimulus statistics. Despite this, the role of adaptive single-neuron nonlinearities in complex tasks remains unclear. In this work, we systematically investigate this question by leveraging recurrent neural networks (RNNs) trained on various tasks.

We introduce a two-parameter family of activation functions that mimics IF-curves of biological neurons, and allows interpolation between well-known nonlinearities used in the modeling and AI literature. Using this novel family in an RNN model, we study the impact of its modulation on: (1) learning dynamics, and (2) activation dynamics, using sequential inference tasks. The proposed parametric family allows for a diversity of dynamical regimes, and we exploit this finding as a prior for selecting fixed nonlinearities, or as a setting for modulation alongside learning. We observe drastic differences in network performance with optima associated with non-trivial biologically realistic combinations of parameters. Finally, we propose a model for online nonlinearity adaptation during input processing, reminiscent of cortical fast neural adaptation. We find that fast adaptation improves performance, and that crucially, it can considerably help mitigate previously unobserved changes in input statistics during a task. We draw parallels between this emergent property of our model and gating mechanisms observed in the brain. Our findings suggest a network-level functional role for the shape and modulation of single-neuron nonlinearities, and reveal non-trivial impacts on the training of artificial neural networks.

1-108. Synaptic metaplasticity in binarized neural networks

Axel Laborieux[1,2] AXEL.LABORIEUX@C2N.UPSACLAY.FR
Maxence Ernoult[3] ERNOULT.M@GMAIL.COM
Tifenn Hirtzlin[1] TIFENN.HIRTZLIN@C2N.UPSACLAY.FR
Damien Querlioz[1] DAMIEN.QUERLIOZ@C2N.UPSACLAY.FR

[1]Paris-Saclay University
[2]Center for Nanoscience and Nanotechnology
[3]Universite de Montreal, Mila

Unlike the brain, artificial neural networks, including state-of-the-art deep neural networks for computer vision, are subject to "catastrophic forgetting": they rapidly forget the previous task when trained on a new one. Neuroscience suggests that biological synapses avoid this issue through the process of synaptic consolidation and metaplasticity: the plasticity itself changes upon repeated synaptic events. In this work, we show that this concept of metaplasticity can be transferred to a particular type of deep neural networks, binarized neural networks (BNNs), to reduce catastrophic forgetting. BNNs were initially developed to allow low-energy consumption implementation of neural networks. In these networks, synaptic weights and activations are constrained to -1,+1 and training is performed using hidden real-valued weights which are discarded at test time. Our first contribution is to draw a parallel between the metaplastic states of and the hidden weights inherent to BNNs. Based on this insight, we propose a simple synaptic consolidation strategy for the hidden weight. We justify it using a tractable binary optimization problem, and we show that our strategy performs almost as well as mainstream machine learning approaches to mitigate catastrophic forgetting, which minimize task-specific loss functions, on the task of learning pixel-permuted versions of the MNIST digit dataset sequentially. Moreover, unlike these techniques, our approach does not require task boundaries, thereby allowing us to explore a new setting where the network learns from a stream of data. When trained on data streams from Fashion MNIST or CIFAR-10, our metaplastic BNN outperforms a standard BNN and closely matches the accuracy of the network trained on the whole dataset. These results suggest that BNNs are more than a low precision version of full precision networks and highlight the benefits of the synergy between neuroscience and deep learning.

1-109. Monkey can play Pac-Man with dynamical hybrid strategy

Qianli Yang[1] QLYANG@ION.AC.CN
Jiaqi Zhang[2] JIAQI_ZHANG2@BROWN.EDU
Jianshu Li[1] JSLI@ION.AC.CN
Wenyi Zhang[1] WYZHANG@ION.AC.CN
Ce Ma[3] CM4913@NYU.EDU
Tianming Yang[1] TYANG@ION.AC.CN

[1]Institute of Neuroscience, Chinese Academy of Sciences
[2]Department of Computer Science, Brown University
[3]New York University Shanghai

The study of the neural mechanisms of animal's higher cognitive functions has been impeded by relatively simplistic behavior paradigms. A behavior paradigm that is well-controlled yet complicated enough to encompass multiple dimensions of cognitive processes will help us understand why a massive computing resource is devoted to cognition in the primates. To address this issue, we trained two macaque monkeys to play a faithful reproduction of the classic video game Pac-Man, in which they used a joystick to navigate the PacMan in a maze where they collected rewards and avoided being eaten by any ghosts. The monkeys understood the key concepts of the games. They were able to clear a maze within on average 5 attempts. Not only did they recognize different states of the ghosts and escaped or approached them accordingly, they also learned to use energizers to turn the ghosts into an edible mode to their advantages. Their decision-making could be characterized as containing a local and a global component, which trade off immediate rewards with a long-term payoff. To further understand the decision-making patterns in the monkeys' game-play, we designed both a rule-based labeling method and a machine learning (ML) multi-agent algorithm to characterize their decision making into different high-level strategies, including a local strategy, a global strategy, and several additional strategies dealing with the ghosts. These two labeling approaches evolved from different perspectives can validate against each other and provide consistent interpretations about monkeys' dynamical internal motivations. In particular, the ML multi-agent algorithm trained with the monkeys' behavior achieved over 95% prediction accuracy on monkeys' joystick movement. Our study provides a behavioral paradigm to study animal cognition in the laboratory setting that is both cognitively sophisticated and fully quantifiable. Future investigations of the neural activities during monkeys' game-playing will have the potential to reveal the neural mechanism of cognition at a complexity level that has been impossible in the field.

1-110. An emergent neural coactivity code for dynamic memory

Mohamady El-Gaby[1,2] MOHAMADY.EL-GABY@NDCN.OX.AC.UK
Vitor Lopes Dos Santos[3] VITOR.LOPESDOSSANTOS@BNDU.OX.AC.UK
David Dupret[3] DAVID.DUPRET@BNDU.OX.AC.UK
Hayley Reeve[4] HAYLEYMREEVE@GMAIL.COM
Natalia Campo-Urriza[1] NATALIA.CAMPO-URRIZA@PHARM.OX.AC.UK
Pavel Perestenko[1] PAVEL.PERESTENKO@PHARM.OX.AC.UK
Alexander Morley[1] ALEX0MORLEY@GMAIL.COM
Lauren Strickland[1] LAUREN.STRICKLAND@PHARM.OX.AC.UK
Istvan Lukacs[1] ISTUAN.LUKACS@BNC.OX.AC.UK
Ole Paulsen[5] OP210@CAM.AC.UK

[1]University of Oxford
[2]Nuffield Department of Clinical Neurosciences
[3]MRC BNDU, University of Oxford
[4]University of Bristol
[5]University of Cambridge

Neural correlates of external variables provide potential internal codes that guide an animal's behaviour. Notably, first-order features of neural activity, such as single-neuron firing rates, have been implicated in encoding information. However, the extent to which higher-order features, such as multi-neuron coactivity, play primary roles in encoding information or secondary roles in supporting single-neuron codes remains unclear. Here we show that millisecond-timescale coactivity between hippocampal CA1 neurons encodes short-lived behavioural contingencies that mice must rapidly learn and later dynamically retrieve. This contingency discrimination was unrelated to the tuning of individual neurons but instead an emergent property of their short-timescale coactivity. Moreover, contingency discriminating coactivity developed with learning, was reactivated selectively during post-learning sharp-wave ripples and was reinstated in a memory probe test in a manner that predicted trial-by-trial performance. Finally, optogenetic suppression of inputs from the upstream CA3 region selectively during learning impaired coactivity-based contingency information in CA1 and subsequent dynamic memory retrieval. These findings show that short-timescale coactivity can combine neurons irrespective of their individual tuning to collectively represent behaviourally-relevant variables for dynamic memory retrieval.

1-111. Nonparametric Inference of Neural Correlations from Sequential Recordings

Francesco Fumarola[1,2] FRANCESCO.FUMAROLA@RIKEN.JP
Amin Nejatbakhsh[3] MN2822@CUMC.COLUMBIA.EDU
Taro Toyoizumi[4] TARO.TOYOIZUMI@RIKEN.JP

[1]RIKEN Center for Brain Science
[2]Center For Brain Science
[3]Columbia University - Zuckerman Institute
[4]RIKEN

The correlation of activity between cells in a neural circuit is a fundamental property of neural dynamics that by definition must be computed on data from cells observed within the same time window. In neuroscience, however, simultaneous measurements of the variables of interest are often not available or must be obtained at the expense of imaging resolution. We show that in such scenario, recordings can be designed to infer correlations between the activities of pairs of recorded channels by using, in each case, information from a third channel in the network simultaneously observed with both. We discuss how delay-embedding principles can be used to solve the problem and propose a procedure to select the "cue cells" to pivot on for reconstruction. We test our proposed technique on synthetic data as well as calcium imaging, electrophysiology, and cultured neurons. Our method is nonparametric, suitable for online deployment, and we find that it outperforms existing approaches for reconstructing neural correlations. It also supports a specific recording paradigm for large ensembles, which consists in using a restricted field of view and moving it forward at time intervals suitable to locating cue cells and saturating the performance of delay-embedding reconstruction from local activity. We test the recording scheme in simplified form by running a synthetic experiment on spiking networks of geometrically distributed neuron clusters and scanning each overlap for cue cells. Since neural decoding is enhanced by fully accounting for correlations, the technique is recommended when acquiring data for decoding purposes, but can be more in general instrumental to stitching neural datasets in the presence of limited observations.

1-112. Increasing neural network robustness improves match to V1 eigenspectrum and improves V1 predictivity

Nathan Kong NCLKONG@STANFORD.EDU
Justin Gardner JLG@STANFORD.EDU
Anthony Norcia AMNORCIA@STANFORD.EDU

Stanford University

Convolutional neural networks (CNNs) show striking similarities to the ventral visual stream. However, their image classification behaviour on adversarially perturbed images diverges from that of humans. In particular, human-imperceptible image perturbations can cause a CNN to misclassify the image. Recent work suggests that the degree to which a system is robust to these perturbations could be related to the power law exponent, α, of the eigenspectrum of its set of neural responses. Informally, the theory states that if $\alpha<1$, then small perturbations to a stimulus could result in unbounded changes to the neural responses. Here, we test this hypothesis by comparing predictions from a set of standard and "robust" CNNs with neural responses in rodent and macaque primary visual cortex (V1). Specifically, we relate these models' V1 fit quality with their accuracy on adversarial images, and their power law exponents.

For both macaque and rodent neural responses, we found that model correspondence to V1 was correlated with adversarial accuracy. We then investigated the relationship between α and adversarial accuracy. When comparing a non-robust model with its robust counterpart, we found that all robust counterparts had higher α. Between α and V1 response predictivity, we similarly found that the robust counterparts of non-robust models had higher α and higher V1 predictivity. Across architectures, however, there was no relationship between these two quantities. Finally, we showed that several robust optimization algorithms lead to higher α, increased robustness and better V1 correspondence. Moreover, an unsupervised algorithm did not improve these three quantities.

These observations suggest that developing biologically plausible techniques to increase α (e.g., reduce representation dimensionality) for a given architecture may improve tolerance to perturbations and V1 response predictivity. We also suggest that α may be a useful parameter to look at when benchmarking models of the visual system.

1-113. Large-scale benchmarking of deep neural network models in mouse visual cortex reveals patterns similar to those observed in macaque visual cortex

Colin Conwell[1,2] CONWELL@G.HARVARD.EDU
David Mayo[3,4] DMAYO2@MIT.EDU
Michael Buice[5] MICHAELBU@ALLENINSTITUTE.ORG
Boris Katz[3] BORIS@CSAIL.MIT.EDU
George Alvarez[1] ALVAREZ@WJH.HARVARD.EDU
Andrei Barbu[3] ANDREI@0XAB.COM

[1]Harvard University
[2]Psychology
[3]Massachusetts Institute of Technology
[4]CSAIL
[5]Allen Institute

What is the representational structure of mouse visual cortex and how is it shaped? Mice obviously interact with the world and recognize objects but unlike in primates, a majority of research to date suggests the activity of their visual cortex may not be so well described by deep neural networks trained for object recognition. Using the Allen Brain Observatory's 2-photon calcium-imaging dataset of activity in over 30,000 rodent visual cortical neurons recorded in response to natural scenes, we work to resolve this discrepancy and demonstrate that modern neural networks can indeed be used to explain activity in the mouse visual cortex to a more reasonable degree than previously suggested. In so doing, we elucidate at large scale the properties of networks which best match the biological visual system, with both representational similarity analysis and encoding models coming to mostly the same conclusions. Our analysis of 30 object recognition architectures (both pretrained and randomly initialized) from the PyTorch model zoo demonstrates that deeper, thinner residual networks with bypass connections, fewer parameters shared across many convolutions, and higher scores on the ImageNet image-recognition challenge tend to be more predictive of the neural activations in our sample. Additionally, we find a significant degree of overlap between the models that best predict macaque visual cortex (as catalogued by brain-score.org) and those that best predict mouse visual cortex. In concert, these findings help to bolster the mouse brain as a viable source of data for the methods that have been successful thus far in the study of monkey brains, and provide a preliminary set of design targets for building models that can better take advantage of the unparalleled scale, quality, and resolution of data afforded by calcium-imaging in the mouse brain.

1-114. Spatial working-memory constraints on computing and maintaining a decision variable

Kyra Schapiro[1,2] KASCHAPIRO@AOL.COM
Joshua Gold[1] JIGOLD@PENNMEDICINE.UPENN.EDU
Zachary Kilpatrick[3,4] ZPKILPAT@COLORADO.EDU
Kresimir Josic[5] JOSIC@MATH.UH.EDU

[1]University of Pennsylvania
[2]Neuroscience
[3]University of Colorado Boulder
[4]Applied Mathematics
[5]University of Houston

Decisions based on evidence accumulated over time require the brain to maintain a representation of the computed decision variable (the accumulated evidence) in working memory (WM). However, if and how WM limitations affect this process is not well understood, because: 1) these limitations have typically been studied with respect to observed, not computed, quantities; and 2) previous studies examined these limitations with binary choice tasks that lacked the ability to examine trial-by-trial variability of analog variables[1]. Here we examined the temporal degradation of a computed analog decision variable represented in WM under different task conditions. The variance of error in WM reports of analog spatial variables has been shown to be consistent with bump attractor dynamics in neural networks[2], so we modeled the remembered values measured here as a diffusing Brownian particle, as has been done previously [3]. The model was validated against data from human subjects reporting either the observed location of a visual target or the estimated mean location of multiple targets after different delays, under several conditions: 1) different numbers of targets (2 or 5) , 2) different report instructions (target or mean indicated at the beginning or end of the trial), and 3) non-simultaneous target presentation (i.e. sequential evidence accumulation). We tested two primary alternative storage hypotheses: 1) a subject calculated the mean before storage in WM and their estimate diffuses exactly as an observed location during the delay period, or 2) a subject held all constituent locations in WM and averaged only once a decision is needed. Preliminary results suggest that: 1) computed decision variables degrade predictably over time in accordance with WM limitations; and 2) this degradation depends strongly on task conditions and individual differences. These results provide important constraints on future studies of how the brain computes and maintains temporally dynamic decision variables.

1-115. Human susceptibility to subtle adversarial image manipulations with unlimited exposure time

Vijay Veerabadran[1,2] VVEERABA@UCSD.EDU
Michael Mozer[3] MCMOZER@GOOGLE.COM
Jonathon Shlens[4] SHLENS@GOOGLE.COM
Jascha Sohl-Dickstein[4] JASCHASD@GOOGLE.COM
Gamaleldin F. Elsayed[4] GAMALELDIN@GOOGLE.COM

[1]Google Research, Brain Team; University of California, San Diego
[2]Google Brain
[3]Google Research, Brain Team; University of Colorado Boulder
[4]Google Research, Brain Team

Artificial neural networks (ANNs) have revolutionized multiple fields, and were initially inspired by models of biological neural networks (BNNs). A growing body of work finds similarities between behaviors and representations in ANNs and BNNs. However, despite these similarities and shared foundations, ANNs exhibit surprising properties and failures that are generally believed not to exist in biological networks. One of the most dramatic of these failures is a susceptibility to adversarial perturbations, where a nearly imperceptible perturbation added to an input can cause an ANN to behave in a dramatically different fashion, for instance mislabeling an initially correctly identified school bus as an ostrich.

Is human vision susceptible to adversarial perturbations? Past work has shown that when images with adversarial perturbations of intermediate magnitude ($\pm$32 of 256 intensity levels) are shown to humans for a short time (~70 ms), human object classification judgements are perturbed in the same direction as ANN. Other work has found that humans can identify adversarial examples with large magnitude perturbations and extended exposure times. Here, we find that human susceptibility to adversarial examples extends beyond these settings. We display images

with small magnitude perturbations (between ± 2 and ± 32 out of 256 intensity levels) for an unlimited exposure time. We find that even at the smallest ± 2 magnitude, these images perturb human judgements of object class in the same direction as an ANN trained to classify images. These results demonstrate that humans exhibit a peculiarity that was once assumed to be specific to machines and are suggestive that perception in ANNs and BNNs is more similar than commonly believed.

1-116. Identifying learning rules from neural network observables

Aran Nayebi ANAYEBI@STANFORD.EDU
Sanjana Srivastava SANJANA2@STANFORD.EDU
Surya Ganguli SGANGULI@STANFORD.EDU
Daniel Yamins YAMINS@STANFORD.EDU

Stanford University

The brain modifies its synaptic strengths during learning in order to better adapt to its environment. However, the underlying plasticity rules that govern learning are unknown. Many proposals have been suggested, including Hebbian mechanisms, explicit error backpropagation, and a variety of alternatives. It is an open question as to what specific experimental measurements would need to be made to determine whether any given learning rule is operative in a real biological system. In this work, we take a "virtual experimental" approach to this problem. Simulating over a thousand idealized neuroscience experiments with artificial neural networks, we generate a large-scale dataset of learning trajectories of aggregate statistics measured in a variety of neural network architectures, loss functions, learning rule hyperparameters, and parameter initializations. We then take a discriminative approach, training linear and simple non-linear classifiers to identify learning rules from features based on these observables. We show that different classes of learning rules can be separated solely on the basis of aggregate statistics of the weights, activations, or instantaneous layer-wise activity changes, and that these results generalize to limited access to the trajectory and held-out architectures and learning curricula. We identify the statistics of each observable that are most relevant for rule identification, finding that statistics from network activities across training are more robust to unit undersampling and measurement noise than those obtained from the synaptic strengths. Our results suggest that activation patterns, available from electrophysiological recordings of post-synaptic activities on the order of several hundred units, frequently measured at wider intervals over the course of learning, may provide a good basis on which to identify learning rules.

1-117. Striatal indirect pathway mediates action switching via modulation of collicular dynamics

Jaeeeon Lee[1,2] JAYLEE@G.HARVARD.EDU
Bernardo Sabatini[1] BERNARDO_SABATINI@HMS.HARVARD.EDU

[1]Harvard University
[2]Neurobiology

Type 2 dopamine receptor-expressing, or indirect pathway striatal projection (iSPNs), neurons comprise one of two major pathways through the basal ganglia, and are a major drug target for treatment of neuropsychiatric disorders. The function of iSPNs is unclear with proposed roles in suppression of unwanted actions and in refining selection actions or their kinematics. Here, we show that iSPNs can simultaneously suppress and facilitate conflicting motor actions in a lateralized licking task. Activation of iSPNs suppresses contraversive while promoting ipsiversive licking, allowing mice to rapidly switch between alternative motor programs. iSPN activity is prokinetic even when mice are not cued to perform an action. Activity in lateral superior colliculus (lSC), a basal ganglia target, is necessary for performing the task and predicts action. Furthermore, iSPN activation suppresses ipsilateral lSC, but surprisingly, excites contralateral lSC. Targeted dimensionality reduction revealed a lower dimensional subspace of SC activity that contained information that trial type. Stimulation of the iSPNs steered the SC neural trajectory towards the ipsiversive direction along the choice axis, without perturbing the activity in other axes. Thus, our results reveal a previously unknown specificity of iSPNs effects on downstream brain regions, including the ability to excite contralateral regions and trigger motor programs. These results suggest a general circuit mechanism for flexible action switching during competitive selection of lateralized actions.

1-118. Spatio-temporal offset between center and surround underlies a bipolar cell motion detector in the retina

Anna Vlasits[1,2] ANNA-LOUISE.VLASITS@UNI-TUEBINGEN.DE
Sarah Strauss[1] SARAH.STRAUSS@UNI-TUEBINGEN.DE
Maria M Korympidou[1] MARIA-MAGDALINI.KORYMPIDOU@CIN.UNI-TUEBINGEN.DE
Yanli Ran[1] YANLI.RAN@CIN.UNI-TUEBINGEN.DE
Katrin Franke[1] KATRIN.FRANKE@CIN.UNI-TUEBINGEN.DE
Timm Schubert[1] TIMM.SCHUBERT@CIN.UNI-TUEBINGEN.DE
Tom Baden[3] T.BADEN@SUSSEX.AC.UK
Philipp Berens[4] PHILIPP.BERENS@UNI-TUEBINGEN.DE
Thomas Euler[1] THOMAS.EULER@CIN.UNI-TUEBINGEN.DE

[1]University of Tuebingen
[2]Institute for Ophthalmic Research/Centre for Integrated Neuroscience
[3]University of Sussex
[4]Institute for Ophthalmic Research, University of Tubingen

Classically, center-surround receptive fields in vision were thought to operate as spatial filters of visual information. However, more complex filters emerge when the center and surround have distinct temporal properties. We studied the role of center-surround receptive fields in motion detection in the mouse retina, where a canonical circuit of bipolar cells (BCs), starburst amacrine cells (SACs), and ganglion cells computes and transmits motion direction information to the brain. The BC contribution to this computation was thought to be non-directional, but the effect of their complex spatiotemporal receptive fields on motion processing has not previously been considered. Surprisingly, we find that the center-surround receptive field of BCs acts as a motion detector, generating a direction signal used by SACs.

In experiments using a fluorescent sensor of glutamate, we directly observe motion direction sensitive BC glutamate release and we show that this direction-selective excitation is strongest for local motion. We further characterized motion sensitivity across BC populations, and find that BCs in distinct synaptic layers of the inner retina are more sensitive to motion than others, but that all BCs possess motion sensing capabilities. Next, using connectomics and functional imaging datasets to design a biophysical model with realistic and type-specific BC inputs, we demonstrate that BCs with intact surround dynamics pass direction-selective excitation to SACs, and that the surround component of BCs' receptive fields is crucial to direction selectivity. To confirm these model findings, we characterized SAC motion preferences experimentally by imaging calcium output signals in their dendrites.

As BCs provide excitation to all amacrine and ganglion cells, the tuning found here may contribute to motion detection in many retinal circuits. Indeed, many stimuli used for studying motion responses should engage this mechanism. Thus, the motion sensing properties of BC receptive fields may have broad implications for visual processing.

1-119. Hippocampal encoding of continual reinforcement learning features

Samia Mohinta[1,2] MOHINTA1234@GMAIL.COM
Martin Dimitrov[1] MARTIN.DIMI97@GMAIL.COM
Dabal Pedamonti[1] DABAL.PEDAMONTI@BRISTOL.AC.UK
Stephane Ciocchi[3] CIOCCHI@PYL.UNIBE.CH
Rui Ponte Costa[1] RUI.COSTA@BRISTOL.AC.UK

[1]University of Bristol
[2]Department of Computer Science
[3]University of Bern

The hippocampus has long been associated with spatial memory and goal-directed spatial navigation. However, how hippocampal neurons encode task relevant features as animals learn multiple tasks is largely unclear. Here we used demixed Principal Component Analysis (dPCA) to analyse recordings from 612 neurons of rats during ego- and allocentric tasks continually interleaved. We find that CA1 hippocampal neurons encode (i) correct and incorrect reward predictions, (ii) task switching information and (iii) eligibility traces. Next, we compare the behaviour and population representations of animals with deep reinforcement learning algorithms, highlighting similarities and differences. The RL model architectures are consistent with the classical hippocampal three-layered structure. A vanilla hippocampal Deep Q-Network (hcDQN) achieves similar average performance when compared to animal learning, but fails to capture behavioural performance during task switching. Hippocampal DQNs when augmented with continual learning algorithms, namely Elastic Weight Consolidation (EWC) and Synaptic Intelligence (SI), provide a better match with animal behavior. Next we tested a hippocampal DQN with a recurrent CA3 layer (hcDRQN), which best captures animal behaviour. This suggests that animals might rely on recurrent networks to learn multiple tasks which is in contrast with the majority of deep learning algorithms.

Next, we used dPCA to analyse the simulated neural activities, which shows that our models can capture some but not all features observed experimentally. Overall, our results provide insights into the role of the hippocampus in multi-task encoding, while putting forward a framework to explicitly compare biological and machine learning during continual learning.

1-120. Simple neural network models exhibit representational drift

David Lipshutz[1,2] DLIPSHUTZ@FLATIRONINSTITUTE.ORG
Shiva Farashahi[3,4] SFARASHAHI@FLATIRONINSTITUTE.ORG
Anirvan M. Sengupta[5] ANIRVANS.PHYSICS@GMAIL.COM
Dmitri Chklovskii[1] DCHKLOVSKII@FLATIRONINSTITUTE.ORG

[1]Flatiron Institute
[2]Center for Computational Neuroscience
[3]Flatiron institute
[4]Center for computational neuroscience
[5]Rutgers University, Flatiron Institute

Neuronal representations in the hippocampus and neocortex gradually change in time despite no apparent change in the environment or behavior. This phenomenon has been termed “representational drift”. While representational drift has been observed repeatedly, the computational role of these changes and the mechanism driving the changes is unknown. We propose that neuronal representations are solutions of principled objectives whose solution sets define continuous manifolds or “solution manifolds”, and representational drift is the process of a neural representation exploring the solution manifold. The exploration of the solution manifold could be due to intrinsic and extrinsic variability and/or feedback signals which preferentially select specific representations.

As a proof of concept, we numerically test for representational drift in two biologically plausible neural network (NN) implementations of learning algorithms that optimize similarity matching objectives. These similarity matching objectives are natural to consider because their solutions are invariant under continuous symmetric transformations, and so the solution sets define continuous manifolds. In each case, we present the NN with inputs from a constant distribution (to represent a constant environment) and show that the neural representation changes in time (i.e., drifts) while at the same time optimizing the similarity matching objective. By demonstrating that representations naturally change in biologically plausible NNs due to an exploration of the solution manifold, we provide a plausible explanation of the mechanism behind representational drift.

2-001. Habits can arise from model-based planning due to latent cause inference

Neil Garrett[1] NEIL.GARRETT@PSY.OX.AC.UK
Sean Allan*[2] SMALLAN@PRINCETON.EDU
Nathaniel Daw[2] NDAW@PRINCETON.EDU

[1]Department of Experimental Psychology, University of Oxford
[2]Princeton University

Summary. Computational models have separately explored two distinct aspects of conditioning. First, findings that instrumental responses are sometimes sensitive – but sometimes insensitive – to reinforcer devaluation have been argued to reflect the use of two learning mechanisms, model-based vs. model-free (Daw et al. 2005). Second, generalization between contexts in classical conditioning has been studied using latent cause inference (LCI) models (Gershman et al. 2010). These theories describe how animals can group experiences into different covert “causes”, depending on their overlap, leading to differential generalization across them. For instance, if acquisition and extinction training are attributed to different causes, extinguished conditioned responses reappear when animals reinstate the acquisition cause. We present a new model combining elements of both lines, to explore the consequences of LCI for instrumental conditioning. The model nests model-based instrumental learning under a Dirichlet-process mixture model for trial clustering. Surprisingly, despite omitting model-free learning, this reproduces key aspects of habits: persistent instrumental responding in extinction for a devalued reinforcer, even when a consumption test verifies the efficacy of aversion conditioning. The model is inspired by and captures the patterns of results from recent experiments (Bouton et al. under review) showing that response sensitivity to reinforcer devaluation depends on the similarity between aversion conditioning and acquisition/test contexts. The model captures value-insensitive responding even using model-based evaluation within each cause due to failure to generalize aversion learning to the latent cause where the instrumental test is inferred to occur. Aversion experience generalizes more readily during the consumption test because of greater feature overlap (e.g., the absence of the lever). Although these results do not rule out a contribution of model-free learning, they point to

the importance of state inference in instrumental learning.

2-002. How the brain makes visual decisions in an ever-evolving world

Cheng Xue[1,2] CHX31@PITT.EDU
Marlene R Cohen[1] COHENM@PITT.EDU
Lily Kramer[1] LEK82@PITT.EDU

[1]University of Pittsburgh
[2]Center for the Neural Basis of Cognition and Department of Neuroscience

Very few decisions in real life are made in a fully-understood, certain, and static environment. Every decision in a natural environment involves two indispensable processes: 1) identifying the relevant problem to solve in the current environment, and 2) correctly solving the problem believed to be relevant. Perceptual decision-making studies with well-controlled behavioral contexts have revealed a lot about the latter process, but typically the former process is deliberately minimized in laboratory studies for simplicity. However, while few neuropsychiatric disorders affect simple perceptual decisions, many affect decision-making under dynamic environments and/or the ability to update belief about the environment based on past experiences. We trained macaque monkeys to perform two perceptual tasks in a controlled but dynamic environment. By recording the monkeys' behavior and the activity of neuronal populations in visual cortex and parietal cortex, we investigated the neuronal mechanisms underlying perception and task belief as the environment dynamically evolved. Contrary to the common intuition that belief updating and perceptual decisions are separable, hierarchical processes, we found behavioral and neuronal evidence that the two processes are inextricably linked. Specifically, through psychophysics, electrophysiology, and normative behavioral modelling, we found that 1) perceptual decisions are more accurate when subjects feel certain about the environment, 2) retaining flexibility between different behavioral contexts comes at a cost of perceptual sensitivity, 3) fluctuations in visual cortical neurons influence the way task belief is updated. These findings suggest interdependent neuronal mechanisms for perceptual decision-making and belief updating; and thus, challenge the notion of a modular brain. Further investigations using our integrated, multi-method approach may have the potential to reshape theories of decision-making in health and disease.

2-003. An optimal spatial code in the hippocampus: environmental and behavioral effects

Johnatan Aljadeff[1,2] ALJADEFF@UCSD.EDU
Tamir Eliav[3] TAMIR.ELIAV@WEIZMANN.AC.IL
Shir Maimon[3] SHIRRACHEL.MAIMON@WEIZMANN.AC.IL
Nachum Ulanovsky[3] NACHUM.ULANOVSKY@WEIZMANN.AC.IL
Misha Tsodyks[3] MISHA@WEIZMANN.AC.IL

[1]University of California, San Diego
[2]Division of Biological Sciences, Section of Neurobiology
[3]Weizmann Institute of Science

Hippocampal place cells are neurons that fire selectively in a specific region ("place field") of the environment, and are thought to support navigation behaviors. Recently it was shown that in bats flying in a very long 200-m tunnel, place cells exhibit a multiscale multifield code, which is qualitatively different from findings in small environments: individual neurons have multiple place-fields, and place-fields of the same neuron differ widely in size. In these experiments, animals reach locomotion speeds much higher than typical in rodents. We developed a theory to understand why the bat hippocampus chooses this surprising encoding of space, and how other species may represent natural environments during motion. To do so, we generalized Shannon's classical theory to apply to the problem of spatial encoding, by taking into account the magnitude of the decoding error; and by accounting for the fact that decoding is done while location is changing. We show that the information content of a place-cell's spikes depends on the relationship between place-field size and velocity, and on the neuron's input-output nonlinearity. We confirmed our theoretical predictions by examining the field-size-velocity relationship across species (bats, rats) in multiple data-sets where both environment size and velocity span two orders of magnitude; and by quantifying input-output relationships through reanalysis of intracellular measurements from mice navigating in virtual reality. Our work suggests that, when faced with the difficult problem of accurately representing a large environment, the hippocampus can adapt place-field statistics – taking into account both movement statistics and biophysical constraints – to form an optimal code for space. We expect that the principles we have discovered will help explain place-cell activity in animals navigating in complex environments, and will motivate new experimental manipulations that will further elucidate mechanisms supporting the formation of the place-cell code.

2-004. Revealing cell types in vivo via dimensionality reduction and graph clustering of spike waveforms

Kenji Lee[1,2] KENJILEE@BU.EDU
Chandramouli Chandrasekaran[1] CCHANDR1@BU.EDU
Hymavathy Balasubramanian[3] HYMAVATHY.BALASUBRAMANIAN@BCCN-BERLIN.DE
Alexandra Tsolias[4,5] ATSOLIAS@BU.EDU
Stephanie Anakwe[1] SANAKWE@BU.EDU
Maria Medalla[1] MMEDALLA@BU.EDU
Krishna Shenoy[1] SHENOY@STANFORD.EDU

[1]Boston University
[2]Psychological and Brain Sciences
[3]Bernstein Center for Computational Neuroscience
[4]Boston University School of Medicine
[5]Anatomy & Neurobiology

Anatomical, physiological, and transcriptomic studies suggest a diverse range of neuronal cell types. However, current in vivo methods—such as clustering on waveform features—only differentiate between broad-(BS) and narrow-spiking (NS) neurons. This gap between 'known' and 'observable' diversity in vivo limits our understanding of how cell types shape behavior. Here, we developed a new method (WaveMAP) combining non-linear dimensionality reduction (UMAP) with graph clustering on spike waveforms and show that it better reveals candidate cell classes in vivo. We applied WaveMAP to extracellular waveforms recorded with U-probes from macaque dorsal premotor cortex (PMd) during a decision-making task. WaveMAP revealed BS and NS classes are comprised of at least three and five sub-classes of neurons respectively. These sub-classes had distinct physiological, functional, and laminar distribution properties. First, BS neuron sub-classes had low firing rates (FR), late choice-selectivity, and broad laminar distributions concentrated in middle layers—hallmarks of excitatory pyramidal cells. Second, two sub-classes of NS neurons had high FR, early choice-selectivity, and strong decision-related responses. The laminar distribution of these neurons was consistent with our layer-specific histological counts of calbindin-/calretinin- and parvalbumin-positive inhibitory interneuron densities respectively. Third, an NS sub-class had identical FR and functional properties as candidate excitatory neurons—consistent with findings that PMd excitatory neurons possess biophysical machinery capable of producing narrow spikes. Fourth, another NS sub-class had FR, functional properties, and triphasic waveforms consistent with excitatory axons. Fifth, WaveMAP sub-classes explained heterogeneity in decision-related properties of PMd neurons over and above previously reported laminar differences. Notably, candidate cell types were unidentifiable with traditional feature-based methods—WaveMAP produced more candidate cellular sub-classes and these sub-classes were also better separable in terms of both waveform shape and properties. In summary, WaveMAP provides previously unavailable access to candidate cell types in vivo and enables investigation of microcircuit dynamics during behavior.

2-005. Primate PPC neurons track passage of time and orchestrate for successful temporal order memory

Shuzhen Zuo[1,2] SUSANZUO.PSY@GMAIL.COM
Lei Wang[1,2] LEI.WANG.PSY@GMAIL.COM
Zhiyong Jin[1] 961292485@QQ.COM
Makoto Kusunoki[3] MAKOTO.KUSUNOKI@PSY.OX.AC.UK
Yong-di Zhou[4] YDZHOU.ICN@GMAIL.COM
Sze Chai Kwok[5] SZE-CHAI.KWOK@ST-HUGHS.OXON.ORG

[1]East China Normal University
[2]School of Psychology and Cognitive Science
[3]Department of Experimental Psychology, University of Oxford
[4]School of Psychology, Shenzhen University
[5]Division of Natural and Applied Sciences, Duke Kunshan University

Despite its relevance in episodic memory retrieval, the role of posterior parietal cortex (PPC) in processing and retrieving temporal information is not completely clear. To explore this, we recorded 676 neurons from the PPC of two monkeys with 32 microelectrodes while they performed a temporal order judgement (TOJ) task between two frames extracted from naturalistic videos. For the video-watching stage, our linear discriminant analysis (LDA) decoder can reliably estimate the passage of time over the time course of a video. The decoding error from the cross-validated LDA was lower than the decoding error from training with 1000 permutations (P<0.001), suggesting that the population of PPC neurons code for the passage of time during video-viewing. Concerning the TOJ period, we compared the activities between correct and incorrect trials. The response of 36 neurons

differentiated between correct and incorrect retrieval. We also examined the pattern similarity in the neuron population between video-encoding and TOJ period to compare between correct and incorrect trials and found that correct trials showed greater similarity between encoding and TOJ than incorrect trials. We further hypothesized that changes of instantaneous synchrony elicited by the neuron populations in its time series would track how memory decisions are made. We used SPIKE-distance to calculate synchronization and found that irrespective of memory outcome, the synchrony level increased and peaked at around 200 ms after onset of the TOJ images. Moreover, the synchrony level in correct trials is significantly, and persistently, higher than incorrect trials before the monkeys made a decision, implying that these neurons work in synchronization to accumulate information for successful TOJ decision. We conclude that the primate PPC neurons carry information about the temporal unfolding of an event and work in concert to support subsequent temporal order memory.

2-006. Spatial differentiation in area LIP dissociated from evidence accumulation

Emilio Salinas[1,2] ESALINAS@WAKEHEALTH.EDU
Joshua Seideman[1] SEIDEMANJA@GMAIL.COM
Terrence Stanford University[1] STANFORD@WAKEHEALTH.EDU

[1]Wake Forest School of Medicine
[2]Neurobiology & Anatomy

The lateral intraparietal area (LIP) contributes to the deployment of spatial attention and planning of eye movements, and also performs sophisticated perceptual operations along spatial and feature dimensions. One such operation is the accumulation of sensory evidence, which is considered an obligatory antecedent to perceptually guided choices. Numerous single-neuron studies in monkeys, typically based on the random-dot motion (RDM) discrimination task, have equated the spatial differentiation of LIP activity with this evidence accumulation process. However, according to recent inactivation experiments, this signal is not necessary for accurate task performance. We propose a simple yet potentially far-reaching explanation for this notable discrepancy: the perceptual evaluation of the motion information occurs more rapidly (~200 ms) than is generally assumed and *precedes* the LIP differentiation. We tested this hypothesis with an urgent version of the RDM task in which there is no systematic lag between the perceptual evaluation and the movement reporting it, and such that the evolution of the subject's choice can be tracked moment by moment. We found that choice accuracy increased steeply with increasing sensory evidence (motion viewing time, in our case). However, at the same time, the spatial signal in LIP became progressively weaker, as if it hindered performance. Analyses of psychometric performance curves conditioned on neuronal activity confirmed this negative association. Furthermore, comparison to a different urgent task in which the relevant sensory evidence was found at the choice targets (rather than near fixation) showed that the RDM result was not due to time pressure alone, but rather to the specific demands placed on the spatial allocation of attention. The results suggest that the ramping activity in area LIP traditionally interpreted as evidence accumulation likely corresponds to a slow, post-decision shift of spatial attention from one location (where the dots are) to another (where the chosen target is).

2-007. Inverse graphics explains population responses in body-selective regions of the IT cortex

Hakan Yilmaz[1,2] HAKAN.YILMAZ@YALE.EDU
Aalap D Shah[1] AALAP.SHAH@YALE.EDU
Ariadne Letrou[1] ARIADNE.LETROU@YALE.EDU
Satwant Kumar[3] SATWANT.KUMAR@UTEXAS.EDU
Rufin Vogels[4] RUFIN.VOGELS@KULEUVEN.BE
Joshua B Tenenbaum[5] JBT@MIT.EDU
Ilker Yildirim[1,6] ILKER.YILDIRIM@YALE.EDU

[1]Yale University
[2]Department of Psychology
[3]The University of Texas at Austin
[4]KU Leuven
[5]Massachusetts Institute of Technology
[6]Psychology

Vision does not merely detect and recognize patterns and contours, but makes rich inferences about objects and agents including their three-dimensional (3D) shapes and configurations. Current modeling approaches based on DCNNs aim to progressively “untangle” category information through non-linear feature hierarchies making image

class labels linearly decodable. These models can explain aspects of the variance in neural responses, but they do not address how perception can be so rich and they typically do not provide a “cognitive lens”, i.e., an interpretable functional account of neural computation. To address these shortcomings, we take a different approach based on “efficient inverse graphics”, the hypothesis that the goal of visual processing is to efficiently invert generative models of how 3D scenes form and project to images. Instead of classification, we use DCNNs to build inference networks, providing mechanistic proposals for how the brain might invert generative scene models. These models meet the functional goal of quickly computing rich percepts and enable a reverse-engineering account of neural computation in the language of objects and generative models. We tested this approach in body perception, where instantaneously interpreting other’s body shapes and configurations is key to flexible behavior. We used an inference network to recover body shape and posture from single images in an articulated generative model. We compared layer activations of this model to population-level activity measured via single-cell recordings in the body-selective regions of macaque IT. Population activity was well explained by inference network layers, including those that map to the stages in the generative model. In contrast, DCNNs trained for classification failed to reproduce the key patterns observed in the neural data. These results contribute to a systems-level understanding of how the brain transforms sense inputs into objects and agents, into rich discrete structures that we can plan with and think about.

2-008. Emergence of sensory and memory representations in the hippocampus and prefrontal cortex of mice

Bryan Souza[1] B.DACOSTASOUZA@DONDERS.RU.NL
Jan Klee[2] J.KLEE@DONDERS.RU.NL
Luca Mazzucato[3] LMAZZUCA@UOREGON.EDU
Francesco Battaglia[2] F.BATTAGLIA@DONDERS.RU.NL

[1]Donders Institute for Brain, Cognition and Behaviour, Radboud University
[2]Donders Institute for Brain, Cognition and Behaviour, Radboud University, Netherlands
[3]Institute of Neuroscience, University of Oregon, USA

The hippocampus and the medial prefrontal cortex (mPFC) are known to participate in episodic and working memory. In trace conditioning, animals learn associating a sensory stimulus (CS) and a reward (US) separated by a delay (trace period). Although the neural activity in both areas can be linked to memory representation formation, it is unknown how their coordinated neural activity may maintain the stimulus-action map evoked during the CS until behavior is expressed. To address this, we use auditory trace conditioning in head-fixed mice with simultaneous 128-channel silicon probe recordings from the hippocampal CA1 subfield and the medial PFC. First, we found that both CA1 and mPFC encoded time within the trial via the emergence of neuronal sequences at the beginning of stimulus presentation. Interestingly, the presence of this time-structured firing was restricted to correct trials (i.e., anticipatory or non-anticipatory behavior in CS+ or CS- trials, respectively) in overtrained sessions. We also found that the representation of the stimulus identity was strong and stable in mPFC, emerging with the stimulus onset, while it was weaker in CA1, mildly increasing during the trace period. Nevertheless, the stimulus decoded from CA1 and mPFC simultaneous activities during trace was strongly correlated after learning, specifically in correct trials. Finally, we found that trial type information (CS+ vs. CS-) could be decoded from neural assemblies transiently active during the trials, and that assemblies from each area followed different reactivation dynamics during sharp-wave ripples happening between trials. Together those results clarify the dynamics of stimulus encoding in the mPFC-hippocampal circuit, suggesting the hippocampus is important to sustain the CS representation during trace period (i.e., in the absence of the stimulus), while mPFC might be needed for creating the stimulus-reward association.

2-009. Neurogym: An open resource to developing and sharing neuroscience tasks

Guangyu Robert Yang[1] GYYANG.NEURO@GMAIL.COM
Jordi Pastor-Ciurana[2] PASTORC.JORDI@GMAIL.COM
Marta Fradera[2] MARTA.FRADERA@GMAIL.COM
Ru-Yuan Zhang[3] RUYUANZHANG@GMAIL.COM
Jeremy Forest[4] JEREM.FOREST@GMAIL.COM
Jorge del Pozo[2] JORGEDELPOZOLERIDA@GMAIL.COM
Joao Barbosa[5,6] PALERMA@GMAIL.COM
Li Ji-An[7] JIAN.LI.ACAD@GMAIL.COM
Christopher J Cueva[1] CCUEVA@GMAIL.COM
Albert Compte[2] ACOMPTE@CLINIC.CAT
Jaime de la Rocha[2] JROCHAV@CLINIC.CAT
Manuel Molano-Mazon[2,8] MANUELMOLANOMAZON@GMAIL.COM

[1]Columbia University
[2]IDIBAPS
[3]Shanghai Jiao Tong University
[4]New York University
[5]Ecole Normale Superieure
[6]Group for Neural Theory
[7]University of California, San Diego
[8]Brain Circuits and Behavior lab

The evaluation of models in computational neuroscience would be greatly facilitated if a given model could be easily trained on a wide range of relevant tasks and compared to many experimental datasets through a variety of analysis methods. However, a major obstacle is the lack of a truly open exchange and easy access to tasks, datasets, and analyses. Towards such an open ecosystem, we are developing Neurogym, a community-driven collection of many neuroscience tasks. The main goal of Neurogym is to provide (1) a large number of modular and customizable tasks written in high-level readable code and with a common Python interface, and (2) ample support for subsequent model development and evaluation. Neurogym builds upon OpenAI Gym, the quintessential collection of reinforcement learning environments, and introduces several key features: (1) A high-level task constructor that allows users to easily build their own tasks and adapt existing ones; (2) A variety of tools called wrappers that allow users to easily assemble complex tasks using simpler ones as modules. (3) Support for training models using either Reinforcement (RL) or Supervised Learning (SL). One key objective of Neurogym is to facilitate the development of neural network models to be trained across many tasks and compared to experimental data. To achieve this goal, on the Neurogym website, we provide tutorial-style codes for training networks on each task. While Neurogym focuses on tasks, to support comparison with experimental data, we open sourced another codebase for efficiently comparing networks with data through a variety of analysis methods. Neurogym is maintained by a growing community of developers and users across the world. The core developers pledge long-term support to Neurogym in the hope of making it a dependable cornerstone of an open ecosystem of computational and systems neuroscience tools.

2-010. Removing independent noise in systems neuroscience data using DeepInterpolation

Jerome Lecoq[1] JLECOQ@GMAIL.COM
Michael Oliver[1] MICHAEL.D.OLIVER@GMAIL.COM
Joshua Siegle[2,3] JOSHS@ALLENINSTITUTE.ORG
Natalia Orlova[1] NATALIAO@ALLENINSTITUTE.ORG
Christof Koch[1] CHRISTOFK@ALLENINSTITUTE.ORG

[1]Allen Institute for Brain Science
[2]Allen Institute
[3]Neural Circuits & Behavior

Progress in nearly every scientific discipline is hindered by the presence of independent noise in spatiotemporally structured datasets. Three widespread technologies for measuring neural activity—calcium imaging, extracellular electrophysiology, and fMRI—all operate in domains in which shot noise and/or thermal noise deteriorate the quality of measured physiological signals. Current denoising approaches sacrifice spatial and/or temporal resolution to increase the Signal-to-Noise Ratio of weak neuronal events, leading to missed opportunities for scientific discovery. Here, we present the development and characterization of DeepInterpolation to remove independent noise from three widely used data modalities in systems neuroscience. Importantly, this algorithm is trained only

on noisy samples from the original raw data, without the need for ground truth information. We characterize its impact on existing data analysis workflows in two-photon calcium imaging, electrophysiological recordings and fMRI imaging. Applying DeepInterpolation to in vivo two-photon Ca2+ imaging yields up to 6 times more segmented neuronal segments with a 15 fold increase in single pixel SNR, uncovering network dynamics at the single-trial level. In extracellular electrophysiology recordings, DeepInterpolation recovered 25% more high-quality spiking units compared to a standard data analysis pipeline. On fMRI datasets, DeepInterpolation increased the SNR of individual voxels 1.6-fold. All these improvements were attained without sacrificing spatial or temporal resolution. DeepInterpolation could well have a similar impact in other domains for which independent noise is present in experimental data.

2-011. Building a generalizable brain-computer interface via fast exploration of decoder parameter space

Sudarshan Sekhar SUDSA89@GMAIL.COM
Emily Oby EMO22@PITT.EDU
Nicole McClain NTM10@PITT.EDU
Aaron Batista APB10@PITT.EDU
Patrick Loughlin LOUGHLIN@PITT.EDU

University of Pittsburgh

Brain computer interfaces (BCIs) have shown promise in restoring movement function to patients with motor dysfunction. Clinically viable BCIs will need to work well on a variety of tasks. Currently, BCIs perform best on the task for which they are calibrated, and it is desirable to create a system that can enable users to perform a variety of tasks. We designed a two-fold strategy to build a generalizable decoder: (1) We developed an adaptive decoder that modifies the parameters of a common BCI decode algorithm, the Kalman filter, to reduce the error between decoded and intended movement. Importantly, the loss function of our adaptive decoder depends only on the observed neural firing rate. This is significant because adaptive decoders typically optimize a loss function that depends explicitly on the kinematics, which is not practical for a real-time BCI application. (2) As we adapt the Kalman parameters to minimize decoding error, we learn the association between neural activity and the space of Kalman parameters generated by the adaptive decoder. We model this association using nonlinear methods such as cubic splines. Finally, we merge the spline fit with the Kalman filter to create a hybrid “SplineKalman” decoder. The SplineKalman continually updates its parameters, which means that it can provide for seamless adjustments of the Kalman filter whenever the task changes. Our novel hybrid decoder outperforms traditional decoders at generalization, and it also outperformed a deep neural network-based BCI decoder.

2-012. Topographic organization in deep neural networks: From network topology to visuotopic organization

Dina Obeid[1,2] DINAOBEID@SEAS.HARVARD.EDU
Talia Konkle[1] TALIA_KONKLE@HARVARD.EDU

[1]Harvard University
[2]John A. Paulson School Of Engineering And Applied Sciences, Department of Computer Science

Neurons in the visual cortex have a consistent topographic organization, where nearby neurons in the cortex have nearby receptive fields, forming a map of the visual field. This topographic map, known as a visuotopic map, is repeated multiple times throughout the visual cortex, with each repetition corresponding to a different visual area. In this work we explore whether such visuotopic organization can emerge as a result of minimizing the total wire length between neurons connected in a deep hierarchical network. In particular we ask, given N neurons with a given connectivity and a 2-d grid with N locations, how will the neurons be placed on the grid such that the total distance between the connected neurons is minimized? This problem is an NP-hard combinatorial problem which we solve using simulated annealing. We first construct deep feedforward hierarchical networks, and examine how different parameters of the network such as filter size, number of channels, and stride affect the placement of the neurons on the grid. By introducing visual input to the network we can visualize the resulting visuotopic organization on the 2-d grid. Our results show that networks with purely feedforward connectivity typically result in a single visuotopic map, and in certain cases no visuotopic map emerges. However, when we modify the network by introducing lateral connections, with sufficient lateral connectivity among neurons within layers, multiple visuotopic maps emerge, where some connectivity motifs yield mirrored alternations of visuotopic maps–a signature of biological visual system areas. These results demonstrate that different connectivity profiles have different emergent organizations under the minimum wiring hypothesis, and highlight that characterizing the large-scale spatial organizing of tuning properties in a biological system might also provide insights into the

underlying connectivity.

2-013. Power-law spatiotemporal scaling in large neural populations follows from widespread coupling to latent dynamical variables

Audrey Sederberg[1] AUDREY.SEDERBERG@GMAIL.COM
Mia Morrell[2] MMORRELL@LANL.GOV
Ilya Nemenman[1] ILYA.NEMENMAN@EMORY.EDU

[1]Emory University
[2]Los Alamos National Laboratory

A key problem in modern neuroscience is extracting useful knowledge from the massive data sets enabled by high-throughput recording technology, such as two-photon imaging of neural populations. Understanding the activity of large populations of neurons is difficult due to the combinatorial complexity of possible cell-cell interactions. Finding a small number of collective variables that explain structure in large-scale activity patterns would be a large step to achieving such understanding. Recently, an averaging procedure, known in the physics community as coarse-graining, was applied to experimental neural recordings from the mouse hippocampus, and this analysis showed over two decades of scaling in free energy, activity variance, eigenvalue spectra, and correlation time, hinting that the mouse hippocampus operates in a critical regime. Operating in a critical regime often means that large-scale collective variables that summarize the activity exist, but it is unclear what mechanisms could explain these scaling results. To understand how this scaling arises, we modeled such data by simulating conditionally independent binary neurons coupled to a small number of long-timescale stochastic fields and then replicating the coarse-graining procedure and analysis. This model reproduces the experimentally-observed scalings, suggesting that scaling in neural recordings is indicative of a small number of latent, dynamical variables that couple widely to the population. Such latent dynamical variables may provide the “collective variables” that explain long-distance and long-time correlations in the activity of populations of neurons.

2-014. Connectome-constrained latent variable model of C. elegans chemosensation

Lu Mi[1] LUMI@MIT.EDU
Richard Xu[2] XUR@JANELIA.HHMI.ORG
Sridhama Prakhya[2] PRAKHYAS@JANELIA.HHMI.ORG
Albert Lin[3] ALBERTLIN@G.HARVARD.EDU
Aravinthan Samuel[3] SAMUEL@G.HARVARD.EDU
Srinivas Turaga[2] TURAGAS@JANELIA.HHMI.ORG

[1]Massachusetts Institute of Technology
[2]HHMI Janelia Research Campus
[3]Harvard University, Department of Physics & Center for Brain Science

Measurements of activity and anatomical connectivity in the C. elegans nervous system can guide the development of detailed computational models. However, the single-neuron and synapse dynamics are generally unknown. Furthermore, the connectome does not directly inform the signs and strengths of individual synapses, and the inputs to sensory neurons are incompletely known. To overcome these challenges, we introduce a connectome-constrained latent variable model (CC-LVM) of the voltage dynamics of the entire C. elegans nervous system.

We use stochastic threshold linear dynamics to model the C. elegans neural network, treating these dynamics as soft constraints on the neural activity. The neuronal voltages are latent variables with prior distributions defined by the dynamics. The chemical and electrical synaptic strengths are learned, but constrained by connectivity given by the anatomical connectome. The framework of variational auto-encoders (VAE) is used to train the model and infer the voltages.

We applied the CC-LVM to whole-brain calcium imaging data which captures 170 of the 300 neurons in C. elegans as it responds to chemosensory stimuli. In principle, an accurate model of the nervous system constrained by incomplete activity measurements but complete connectivity measurements can enable accurate predictions of neural activity in neurons which were not recorded. We tested this hypothesis by using the CC-LVM to predict activities of neurons which were experimentally measured, but whose activity was withheld during model training. The CC-LVM predicted neural activity of withheld neurons, significantly better than models unconstrained by the connectome, demonstrating the utility of the connectome even when little is known about the signs and strengths of individual connections. CC-LVMs thus provide a new tool for connectome and activity constrained modeling of

neural circuits.

2-015. Voltage-based inhibitory synaptic plasticity: network regulation and diversity

Victor Pedrosa VICTORPEDROSABC@GMAIL.COM
Claudia Clopath C.CLOPATH@IMPERIAL.AC.UK

Imperial College London

Neural networks are highly heterogeneous while homeostatic mechanisms ensure that this heterogeneity is kept within a physiologically safe range. One of such homeostatic mechanisms, inhibitory synaptic plasticity, has been observed across different brain regions. Computationally, however, inhibitory synaptic plasticity models often lead to a strong suppression of neuronal diversity. Here, we propose a model of inhibitory synaptic plasticity in which synaptic updates depend on presynaptic spike arrival and postsynaptic membrane voltage. Our plasticity rule regulates the network activity by setting a target value for the postsynaptic membrane potential over a long timescale. In a feedforward network, we show that our voltage-dependent inhibitory synaptic plasticity (vISP) model regulates the excitatory/inhibitory ratio while allowing for a broad range of postsynaptic firing rates and thus network diversity. In a feedforward network in which excitatory and inhibitory neurons receive correlated input, our plasticity model allows for the development of co-tuned excitation and inhibition, in agreement with recordings in rat auditory cortex. In recurrent networks, our model supports memory formation and retrieval while allowing for the development of heterogeneous neuronal activity. Finally, we implement our vISP rule in a model of the hippocampal CA1 region whose pyramidal cell excitability differs across cells. This model account for the experimentally observed variability in pyramidal cell features such as the number of place fields, the fields sizes, and the portion of the environment covered by each cell. Importantly, our model supports a combination of sparse and dense coding in the hippocampus. Therefore, our voltage-dependent inhibitory plasticity model accounts for network homeostasis while allowing for diverse neuronal dynamics observed across brain regions.

2-016. Multivariate temporal receptive fields in speech perception reflect low-dimensional dynamics

Emily Stephen EMILYPS14@GMAIL.COM
Yuanning Li YUANNING.LI@UCSF.EDU
Sean Metzger SEAN.METZGER@UCSF.EDU
Yulia Oganian YULIA.OGANIAN@UCSF.EDU
Edward Chang EDWARD.CHANG@UCSF.EDU

University of California, San Francisco

Research into the cortical basis of auditory speech perception has successfully modeled how high gamma (70-150Hz) responses in electrocorticography (ECoG) over the superior temporal gyrus (STG) encode phonetic features such as consonants and vowels. Furthermore, some STG populations respond to sentence onsets and acoustic onset edges (“peak rate” events), which represent sentence-level and syllable-level timing, respectively. This timing information could be important for integrating phonetic feature responses into higher order representations. We asked whether these timing representations are shared across populations in a low-dimensional latent state. If so, it would indicate spatially distributed dynamics that could be used locally to identify the relative timing of phonetic features with respect to the sentence and syllable timescales. Estimating a latent state for individual features is difficult because speech stimuli induce temporally extended responses that overlap in time. Unsupervised clustering can infer overall latent state, but it cannot disentangle responses related to different stimulus features. Hence we use a supervised approach, integrative reduced rank regression (iRRR), which estimates separate latent dynamics for each feature. We show that iRRR outperforms models that treat each electrode individually, while using low-dimensional representations of the states for each feature. This indicates that substantial feature-related information is shared across electrodes. Sentence onset and peak rate events explain more of the variance than phonetic features, and their spatial support shows a clear division between posterior and middle regions of STG: posterior STG has stronger sentence onset responses and shorter-latency peak rate responses than middle STG. Dynamically, the latent states for the sentence onset and peak rate features reveal rotational dynamics that capture continuous relative timing information. We propose that sentence onset and peak rate latent states represent a spatially distributed timing signal that could be used locally, for example when composing higher-order representations from low-level acoustic features.

2-017. Reverse engineering what the retina encodes about natural motion

Siwei Wang[1,2] SIWEIW@UCHICAGO.EDU
Stephanie Palmer[1] SEPALMER@UCHICAGO.EDU

[1]The University of Chicago
[2]Organismal Biology and Anatomy

Recent work has shown that natural images can be reconstructed from the population activity of retinal ganglion cells (Brackbill et al., 2020, Parthasarathy & Simoncelli 2020). However, the behavioral needs of an organism are embedded in a complex dynamical environment. Whether it is to find a prey, search for a mate, or escape from a predator, an animal must extrapolate natural motion trajectories. Does it "reuse" static representations for core vision tasks, i.e., those learned from static scene textures to encode natural motion? To investigate this, we build a generative neural network combining a unet (Ronneberger et al., 2015) and a variational autoencoder (Kingma & Welling 2013) which uses retinal spike trains to reconstruct natural scene frames. Said another way, we use spikes trains generated by past frames (what the retina has seen) to predict future frames (what the retina is going to see). By using "early stopping" (15 epochs total), we ensure that the trained neural architecture encodes general features from the input spike trains. By comparing dynamic and static generative models, we see that motion is best encoded in a lower dimensional feature space than the static scene. We also show that the retina is able to predict spatiotemporal features beyond the timescale of its single neuron autocorrelation (~100ms). By comparing the features learned from dynamic generation with static reconstruction, we observe that the encoding of natural motion highlights that certain spatiotemporal dynamic features are missed in static reconstruction. This may suggest novel coding mechanisms in the sensory population for behaviorally relevant predictive features.

2-018. A neural signature of anticipation in the macaque frontal cortex

Nicolas Meirhaeghe[1,2] NMRGHE@MIT.EDU
Hansem Sohn[1,3] HANSEM@MIT.EDU
Mehrdad Jazayeri[1] MJAZ@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Harvard-MIT Division of Health Sciences & Technology
[3]McGovern Institute for Brain Research

Humans continuously form expectations about the world based on prior experience and use this information to improve their future behavioral responses. For instance, when a batter prepares to hit an incoming baseball, they may adjust their motor plan based on their prediction of the ball's speed estimated from the past few throws. How are these types of temporal expectations encoded in the brain? Here we offer a potentially general answer based on recordings from the dorsomedial frontal cortex (DMFC) of monkeys performing a time interval reproduction task. We found that temporal expectations during the measurement of an interval are encoded by the speed at which neural trajectories evolve. The speed is adjusted such that the population activity crosses a particular state at the precise moment when elapsed time matches the mean of previously encountered intervals. This mechanism resembles the control process that was previously proposed for initiating delayed movements (Wang et al., 2017). However, here, the speed control was not associated with movement initiation; instead, the speed reflected the animals' expectations of an upcoming event based on past experience. This result was supported by a second experiment in which we covertly changed the underlying temporal statistics of the task. After the switch, the speed of neural trajectories changed so as to reflect the new expected interval. Remarkably, this adjustment occurred several hundreds of trials before animals showed any behavioral signs of adaptation. We propose that the function of this speed control mechanism is to enable the brain to measure time intervals directly in terms of their deviation from expectation, i.e., sensory prediction errors.

2-019. Normalized value coding provide a computational mechanism for distributional reinforcement learning

Kenway Louie[1,2] KL837@NYU.EDU

[1]New York University
[2]Center for Neural Science

Novel, uncertain, or dynamic environments require organisms to learn appropriate behavior based on environmental feedback. This learning is widely modeled in psychology, neuroscience, and computer science by prediction error-guided reinforcement learning (RL) algorithms. In standard RL theory, reward predictions are represented

as a single scalar quantity and agents learn the mean value of stochastic outcomes. In contrast, recent theoretical and empirical work suggest that neural systems employing asymmetric prediction errors can instead learn the full distribution of possible rewards; however, the biological basis of this asymmetry is unknown. Here, we characterize the behavior of an RL algorithm with a specific nonlinear value function implemented by divisive normalization and show that normalized RL (NRL) introduces an intrinsic and tunable asymmetry in prediction error coding. At the behavioral level, this asymmetry explains empirical variability in risk preferences typically attributed to asymmetric learning rates. At the neural level, diversity in NRL asymmetries replicates empirical dopaminergic biases and provides a computational mechanism for distributional RL, learning the full probability distribution of future rewards. Together, these results suggest that nonlinear RL algorithms provide advantages in behavioral and computational flexibility, and argue for an incorporation of biologically valid value functions in computational models of learning and decision-making.

2-020. Computational principles of value construction

Kiyohito Iigaya[1] KIIGAYA@CALTECH.EDU
Sanghyun Yi[1] SYI@CALTECH.EDU
Iman Wahle[1] IWAHLE@CALTECH.EDU
Sandy Tanwisuth[2,1] K.TANWISUTH@GMAIL.COM
Aniek Fransen[1] AFRANSEN@CALTECH.EDU
John O'Doherty[1] JDOHERTY@CALTECH.EDU

[1]California Institute of Technology
[2]University of California, Berkeley

It is an open question how humans and other animals flexibly construct the value of stimuli, even for stimuli never encountered before. While great progress has been made toward understanding how the brain adjusts the value of stimuli through reinforcement-learning, little is known about how stimulus value arises in the brain in the first place. Here, we propose and provide evidence that the brain constructs the value of a novel stimulus by extracting and assembling common features. Notably, because those features are shared across a broad range of stimuli, we show that simple linear regression in the feature space can work as a single neural mechanism to construct the value across stimulus domains. In large-scale behavioral experiments with human participants, we show that a simple model of feature abstraction and linear summation can predict the subjective value of paintings, photographs, as well as shopping items whose values change according to different goals. The model shows a remarkable generalization across stimulus types and participants, e.g. when trained on liking ratings for photographs, the model successfully predicts a completely different set of art painting ratings, or when trained on clothing items for one gender, the model successfully predicts ratings in the other gender across contexts. Also, we show that these general features emerge through image recognitions training in a deep convolutional neural network, without explicit training on the features, suggesting that features relevant for value computation arise through natural experience. Furthermore, using fMRI, we found evidence that the brain actually performs value computation hierarchically by transforming low-level visual features into high-level abstract features which in turn are transformed into valuation. We conclude the feature-based value computation is a general neural principle enabling us to make flexible and reliable value computations for a wide range of stimuli.

2-021. Breadth versus depth: identifying close-to-optimal heuristics in human decision-making with finite resources

Alice Vidal[1,2] ALICE.VIDAL@UPF.EDU
Ruben Moreno Bote[1,3] RUBEN.MORENO@UPF.EDU
Salvador Soto-Faraco[4] SALVADOR.SOTO@UPF.EDU

[1]University Pompeu Fabra, Barcelona
[2]Center for Brain and Cognition (CBC)
[3]Center for Brain and Cognition
[4]University Pompeu Fabra

Planning a holiday trip, selecting university courses or investing in stock options, these situations require the allocation of finite resources (time, money) in the absence of immediate feedback. Such cases can be described under the framework of the breadth-depth dilemma; a multi-alternative risk-taking problem where optimality requires striking a delicate balance between pure depth – concentrating all your capacity in one alternative – and pure breadth – spreading your capacity across the most alternatives. Do humans choose optimally in these situations? How do we manage the breadth-depth (BD) trade-off depending on the resources available? Although this question has been previously discussed [1]–[4], only recently it has been formalised through a finite sample

capacity model [5] that makes quantitative predictions defining optimal choice behaviour as a function of available resources (sampling capacity) and likelihood of reward (environment richness). Here, we test the model predictions empirically by manipulating these two variables. As expected, we found that (1) at low capacities pure breadth predominates, whilst at high capacities participants sample just a few alternatives in depth, and intentionally ignore the rest, (2) the number of different alternatives sampled increases with capacity following a power law with an exponent close to ½, and (3) rich environments promote depth over breadth, so that participants tend toward depth as the probability of success increases overall. Interestingly, some divergences between empirical data and predictions also stand out. First, we predominantly observe a gradual transition of the BD trade-off rather than the sharp transition predicted by the model. Second, we identify a bias to sample homogeneously among the selected alternatives, which deviates from optimal BD strategies. To conclude, our study provides a novel framework to investigate multi-alternative decision-making under finite resources, it reveals near-optimal choice behaviour, and identifies some heuristics that simplify computations at the cost of optimality.

2-022. Force cues flexibly separate motor memories in human reaching adaptation

Frederic Crevecoeur[1,2] FREDERIC.CREVECOEUR@UCLOUVAIN.BE
James Mathew[3] JAMES.MATHEW@UCLOUVAIN.BE
Philippe Lefevre[3] PHILIPPE.LEFEVRE@UCLOUVAIN.BE

[1]Universite catholique de Louvain
[2]Institute of Neuroscience
[3]University of Louvain

Current evidence suggests that population dynamics combined with initial neural states associated with movement planning underlines the generation of motor commands1. From this perspective, the ability to learn multiple and potentially opposite force field disturbances applied during reaching may require that preparatory activity be distinct2. Here we tested whether applying a background force prior to a reaching movement could separate motor memories by exposing healthy volunteers to opposite velocity dependent force fields (clockwise and counterclockwise). Our results confirmed that a background force cue about the direction of the upcoming force field could form independent motor memories of the randomly interleaved opposite force fields. Interestingly, this association was flexible, as upon reversing the relationship between the background and force field directions midway through the experiment, participants were able to relearn the association and form novel, independent motor memories of the force field cued by the background force applied in the opposite direction. To check if this associated was sensitive to the magnitude of the cue, we went further and performed a second experiment where we manipulated the magnitude of the force field (only one direction), cued by the magnitude of the background force. We found a partial separation of motor memories: first interference occurred such that participants zeroed the effect of the heavier perturbation, at the cost of overcompensating for the lighter one during the first few trials. Nevertheless, a statistical analysis indicated that over the course of several tens of trials, the anticipation of force field changed independently across the different magnitudes, suggesting that the memories were indeed separated in spite of the interference. To conclude, the formation of motor memories facilitated by force cues of different directions suggests that muscle afferent feedback prior to movement can put preparatory cortical activity into different states, thereby allowing simultaneous formation of independent motor memories.

2-023. Functional mapping of neuronal ensemble connectivity via compressed sensing

Phillip Navarro[1,2] PHILLIPNAVARRO@UFL.EDU
Karim Oweiss[1] KOWEISS@ECE.UFL.EDU

[1]University of Florida
[2]Electrical and Computer Engineering

Mapping connectivity between neurons has been traditionally studied by individually stimulating putative presynaptic cells to evoke a response in their postsynaptic counterparts. This method has a very low yield and is not feasible to use when functional mapping of large networks is needed. Holographic Two photon (2P) optogenetic stimulation of single cells combined with ensemble voltage imaging can improve this approach. However, the extent to which this approach can be scaled without significantly increasing experimental time remains largely unknown. Here, we used an in silico model of recurrently connected excitatory and inhibitory population of Izhikevich model neurons to discover monosynaptic connectivity among an observable subset of neurons responding to modeled optogenetic stimulation. By exploiting the sparseness of the network and ensuring independence of simultaneously stimulated subsets of neurons across trials, we show that the functional mapping of the net-

work can be formulated as a compressive sensing problem. In particular, we were able to map the connectivity between neurons in the observed population by solving an optimization problem that minimizes the number of sources needed to fully reconstruct the evoked responses. We demonstrate that the proposed method requires far fewer number of trials to reach > 95% reconstruction accuracy compared to the standard whole cell patch clamp approach. Reconstruction accuracy remained robust to changes in network and stimulation parameters, maintaining results for random weight initializations, probabilities of connection between neurons and number of neurons simultaneously stimulated per trial. Moreover, the method is efficient as the number of observable neurons and trials are varied, highlighting its scalability to larger size populations and limited experimental time.

2-024. Population codes that enable sample efficient learning

Blake Bordelon[1,2] BLAKE_BORDELON@G.HARVARD.EDU
Cengiz Pehlevan[1] CPEHLEVAN@SEAS.HARVARD.EDU

[1]Harvard University
[2]Applied Mathematics

The ability to learn a pattern from a small number of examples is a central component of intelligence. Understanding how and why the brain can learn in such a sample efficient manner remains an open problem. In this work, we study how the structure of neural population codes influences the number of sampled stimuli required to learn an accurate linear readout. We find that sample efficient learning requires a population level inductive bias determined by the structure of the inner-product kernel, which measures the similarity of neural responses to two different input stimuli. We demonstrate that the eigendecomposition of the kernel provides information about the strength of the inductive bias as well as the compatibility of the population code with certain tasks. We apply our theory of generalization to experimental population responses of mouse V1 neurons, showing how they have a good inductive bias for low frequency orientation discrimination tasks. Although many alternative population codes can possess identical kernels, V1 populations show significantly sparser responses than typical codes with the same inductive biases.

2-025. Third-order interaction as a gating mechanism in neural computation

Yicheng Fei[1,2] FYCISC@GMAIL.COM
Xaq Pitkow[3] XP4@RICE.EDU

[1]Rice University
[2]Physics and Astronopy
[3]Baylor College of Medicine, Rice University

To discern if an apple is ripe, it is important to tell whether the surface is red or is just illuminated by red light. The illuminant changes how the surface reflectance affects the image. Although such gating or modulatory interaction is ubiquitous in our environment and in circuits, no mathematical framework exists for solving these problems in general. For tractability, most probabilistic models have been based on mere pairwise interactions, which cannot express many phenomena in neuroscience. For example, they cannot capture triple synapses, feedback gating, or cooperative effects without complex nonlinear latent pairwise interactions. Yet gating is a crucial, common operation in machine learning, as seen in LSTMs, GRUs, and attention networks. Here we describe a statistical foundation for such operations. Because gating between three interacting variables functions much like a transistor, we call this third-order motif a ‘statistical transistor’. This provides a more universal account of contextual interactions observed in neural activity and previously captured by a Mixture of Gaussian Scale Mixtures (MGSM). The MGSM uses a mixture model to describe the conditional correlation structure and produces flexible divisive normalization among V1 neurons. This effect can also be modeled with our more general third-order model. Our statistical transistors extend naturally to broader classes of problems, including population-scale visual contexts and hierarchical structure. We hope this motif will provide a novel normative generalization of common phenomenological models including divisive normalization.

2-026. Inference as control

Lokesh Boominathan[1,2] LB36@RICE.EDU
Paul Schrater[3] SCHRATER@UMN.EDU
Xaq Pitkow[4] XAQ@RICE.EDU

[1]Rice University
[2]Electrical and Computer Engineering
[3]University of Minnessota
[4]Baylor College of Medicine, Rice University

A critical computation for the brain is to infer the world's latent variables from ambiguous observations. Computational constraints, including metabolic costs and noisy signals, limit the performance of these inferences. Efficient coding is a prominent theory that describes how limited resources can be used best. In one incarnation, this leads to the theory of predictive coding, which posits that predictions are sent along feedback channels to be subtracted from signals at lower cortical areas; only the difference returns to the higher areas along feedforward channels, reducing the cost of sending redundant signals already known to the higher areas. This theory does not, however, account for the costs or noise associated with the feedback. Depending on the costs for sending predictions and the reliability of signals encoding those predictions, we expect different optimal strategies to perform computationally constrained inferences. For example, if the feedback channel is too unreliable and expensive, we hypothesize that it is not worth sending any predictions at all. Here we offer a more general theory of inference that accounts for the costs and reliabilities of the feedback and feedforward channels, and the relative importance of good inferences about the latent world state. We formulate the inference problem as control via message-passing on a graph, maximizing how well an inference tracks a target state while minimizing the message costs. Messages become actions with their own costs to reduce while improving how well an inference tracks a target state. We solve this problem under Linear-Quadratic-Gaussian (LQG) assumptions: Linear dynamics and transformations, Quadratic state and control message costs, and Gaussian noise for the process, observations, and measurements. Our theory enables us to determine the optimal predictions and how are they are integrated into computationally constrained inference.

2-027. Modelling continual learning in humans with Hebbian context gating

Timo Flesch TIMO.FLESCH@PSY.OX.AC.UK
David Nagy NAGY.G.DAVID@WIGNER.MTA.HU
Andrew Saxe ANDREW.SAXE@PSY.OX.AC.UK
Christopher Summerfield CHRISTOPHER.SUMMERFIELD@PSY.OX.AC.UK

Department of Experimental Psychology, University of Oxford

Humans continue to learn across the lifespan, progressively building a repertoire of learned tasks. However, in machine learning research, how to layer new learning on old without provoking catastrophic forgetting remains an unsolved challenge (Parisi et al, 2018). Interleaving training data from multiple tasks allows artificial neural networks to find a joint solution that avoids interference. However, humans strikingly perform worse after interleaved training in task-switching paradigms (Flesch et al, 2018). This discrepancy between humans and machines is likely to limit the utility of deep neural networks as models of human cognitive processes. Here, we propose a novel learning algorithm for continual learning in deep neural networks and demonstrate that it captures the benefit of blocking and the cost of interleaving in humans performing a context-dependent discrimination task (Mante et al 2013). We introduce two biologically plausible algorithmic innovations: (i) neuronal responses are "sluggish" so that context cues continue to impact future decisions (sequential effects) (Foldiak 1991); and (ii) supervised learning is supplemented with a Hebbian step in which co-activated input units become preferentially connected (Oja 1992). Together, these assumptions allow the network to learn continually in blocked conditions, but at the cost of interference in interleaved conditions, mirroring the pattern of data observed in humans. Notably, another proposed solution to continual learning known as Elastic Weight Consolidation (Kirkpatrick et al 2016) fails to strongly partition contexts, and thus fails to prevent interference. The implication is that unlike artificial neural networks, biological systems use temporal context to cluster task knowledge via unsupervised learning, which in turn allows blocked (but not interleaved) supervised learning to rely on partitioned input-output mappings. Our contribution provides new insights into the mechanisms underlying continual learning in humans and its implementation at the circuit level.

2-028. Quantitative deconstruction of behavioral ontogeny

Jesse Marshall[1] JESSE_D_MARSHALL@FAS.HARVARD.EDU
Diego Aldarondo[1] DIEGOALDARONDO@G.HARVARD.EDU
William Wang[1] WLWANG70@GMAIL.COM
Amanda Gellis[1] AJG1GELLIS@GMAIL.COM
Bence Olveczky[1] OLVECZKY@FAS.HARVARD.EDU
Tim Dunn[2] TIMOTHY.DUNN@DUKE.EDU

[1]Harvard University
[2]Duke University

Infants must crawl and cruise for months before they walk. Rats play fight for weeks to build motor and social skills for adulthood. Over development, mammals learn to robustly produce a diversity of motor behaviors through an innate learning curriculum. However we lack a quantitative portrait of this motor development process, as studies have largely been limited to measuring behavior using hand annotation. As a result it is unknown what types of motor experience are important and what impact this experience has on the brain. Understanding motor development and its neural basis could yield new insights into neurodevelopmental diseases and the structure of neural representations. It could also aid in the construction of new artificial motor controllers. Existing artificial controllers lack the flexibility and robustness of mammalian behavior, and developmentally inspired learning curricula could be a solution. We recently developed DANNCE, a tool based on geometric deep learning that enables 3D kinematic tracking in diverse species, with orders of magnitude greater precision and robustness than 2D convolutional networks such as DeepLabCut. Using DANNCE we were able to precisely measure the type and usage of different behaviors over development, as well as their 3D kinematics. We found increases in grooming and rearing behaviors, as well as in overall behavioral complexity over development. We then rigorously characterized the lineage structure of different behaviors as they emerged, finding an expansion of rearing behaviors over time, a necessary first step towards understanding behavioral and neural development. Overall, we introduce a grounded framework for understanding the development of mammalian behavior.

2-029. Cerebellar-like structure improves feedback-learning in recurrent neural networks

Alessandro Barri[1] A.BARRI@UCL.AC.UK
Harsha Gurnani[1,2] HARSHA.GURNANI.16@UCL.AC.UK
R. Angus Silver[1] A.SILVER@UCL.AC.UK

[1]University College London
[2]Department of Neuroscience, Physiology and Pharmacology

Artificial recurrent neural networks (RNNs) trained on behavioural tasks can exhibit activity that strikingly resembles that of neocortical networks [Mante et al. 2013]. Such RNNs can, for example, produce realistic motor outputs [Sussillo et al. 2015] and solve cognitive tasks [Yang et al. 2019]. It is thought that analysis of their internal dynamics can lead to a deeper understanding of the principles by which biological networks operate [Musall et al. 2019]. Often, a single, isolated RNN is used as a proxy for a neocortical circuit. However, neocortex operates in conjunction with other brain areas, including the cerebellum, forming cortico-cerebellar loops [Kelly & Strick 2003]. Classical models of cerebellar function propose that activity patterns are expanded into higher-dimensional representations in the cerebellar granule cell layer (GCL), improving learning [Marr, 1969; Albus, 1971; Litwin-Kumar et al. 2017; Cayco-Gajic et al. 2017]. To investigate whether cerebellar-like expansion recoding could improve learning in neocortical-like RNNs we compared time series learning using the FORCE algorithm [Sussillo & Abbott 2009] in two distinct RNN architectures. The first consisted of a single RNN (S-RNN), which received feedback from its output unit(s), while in the second a cerebellar-like GCL network was connected to the output unit(s) (C-RNN). The C-RNN outperformed the S-RNN in time series learning, regardless of the levels of noise injected into the recurrent units, when the fraction of active units in the GCL ranged from 10%-75%. Moreover, the relative performance of the C-RNN increased with the dimensionality of the target signal. Our results show that expansion recoding can improve multidimensional temporal sequence learning in neocortical-like RNNs, and suggest a defined computational role for cortico-cerebellar loops.

2-030. Controlling robotic locomotion by a closed-loop interaction with living central pattern generators

Rodrigo Amaducci[1] RODRIGO.AMADUCCI@UAM.ES
Irene Elices[2] IRENE.ELICES@INSERM.FR
Manuel Reyes-Sanchez[1,3] MANUEL.REYES@UAM.ES
Alicia Garrido-Pena[1,4] ALICIA.GARRIDO@UAM.ES
Rafael Levi[1] RLEVI66@GMAIL.COM
Francisco B. Rodriguez[1] F.RODRIGUEZ@UAM.ES
Pablo Varona[1,5] PABLO.VARONA@UAM.ES

[1]Universidad Autonoma de Madrid
[2]Institut de la Vision, Sorbonne Universite
[3]Escuela Politecnica Superior
[4]Departamento Ingenieria Informatica. Computational Neuroscience Group
[5]Ingenieria Informatica

A hybrid robot or hybrot is a technological combination of biological circuits with robotics. This technology allows validating the role of specific neural dynamics involved in the control of locomotion, sensorimotor transformation and behavior, and thus it is a step forward towards future hybrid entities. Hybrots are still largely undeveloped and are typically implemented with neuron cultures and multichannel electrode arrays. On the other hand, bio-inspired central pattern generators (CPG) have been widely used in robotic implementations due to their capacity to generate autonomous and robust rhythmic sequences that can coordinate motor functions. However, the richer intrinsic dynamics of living CPG have never been used for hybrid robot locomotion control.

In this work, we present the first hybrot built with a living CPG. Using an advanced dynamic-clamp protocol, the pyloric CPG from a Carcinus maenas is connected to the motor plant of the robot. The dynamical invariants between the intervals that build the CPG sequence are used for the coordination of the robot locomotion. Sensory feedback is injected into the CPG in the form of intracellular current according to the robot's light sensor activation. The connection between the living neurons and the robot is established via Bluetooth following an online closed-loop protocol. We report that, when modulated and coordinated by the living circuit flexible rhythm, an effective robotic locomotion is accomplished and the existing dynamical invariants in the form of cycle-by-cycle linear relationships between the activation intervals of the neurons are transferred into a coherent movement. At the same time, the hybrot generates appropriate locomotor responses to changes in the environment as detected by the robot's sensor.

2-031. Testing the state dependent model of time perception against experimental evidence

Pirathitha Ravichandran-Schmidt[1] PIRATHITHA.RAVICHANDRAN-SCHMIDT@SRH.DE
Joachim Hass[2] JOACHIM.HASS@SRH.DE

[1]SRH University Heidelberg and "Ruprecht Karl University Heidelberg"
[2]SRH University Heidelberg

<p>Coordinated movements, speech and other actions are impossible without precise timing. Realistic computational models of interval timing in the mammalian brain are expected to provide key insights into the underlying mechanisms of timing. Existing computational models of time perception have only been partially replicating experimental observations, such as the linear increase of time, the dopaminergic modulation of this increase, and the scalar property, i.e., the linear increase of the standard deviation of temporal estimates with time (Buhusi and Meck, 2005). In this work, we incorporate the state-dependent computational model, which encodes time in the dynamic evolution of network states without the need for a specific network structure (Buonomano, 2000) into a biologically plausible prefrontal cortex (PFC) model based on <i>in</i> <i>vivo</i> and <i>in</i> <i>vitro</i> recordings of rodents (Hass, Hertag and Durstewitz, 2016). We show that the naturally occurring heterogeneity in cellular and synaptic parameters in the PFC is sufficient to encode time over several hundreds of milliseconds. The readout faithfully represents the duration between two stimuli applied to the superficial layers of the network, thus fulfilling the requirement of a linear encoding of time. A simulated activation of the D2 dopamine receptor leads to an overestimation of time, in line with experimental results (Buhusi and Meck, 2005). Furthermore, we show that the scalar property holds true for intervals of several hundred milliseconds. We conclude that this model is capable of representing durations up to 400 ms in a biophysically plausible setting, compatible with experimental findings in this regime. For larger durations, however, other mechanisms will likely come into play, as has been suggested by a recent theoretical study (Hass and Herrmann, 2012).</p>

2-032. Subjective experience carried by premotor output to striatum supports decision-making

Drew Schreiner[1,2] DCSCHREI@UCSD.EDU
Christina Gremel[1] CGREMEL@UCSD.EDU

[1]University of California, San Diego
[2]Psychology

Subjective experience is a powerful driver of decision-making, but the majority of neurobiological studies ignore its contribution. Decision-making investigations often institute a trial structure, limit choice and movement, and elicit behavior. Neurobiological investigations likewise interpret associated mechanisms without accounting for the subjective, experiential nature of such decision-making. Indeed, the seemingly widespread distribution of similar decision-making information across the brain may in part be due to failing to account for temporal and contextual information accrued by subjective experience that drives decision-making. Here we investigate how subjective experience contributes to decision-making and ask how corticostriatal circuits may differentially represent and control this experience. Mice need to hold down a lever-press for a minimum duration to earn food. All animals have is their prior experience, as performance is self-initiated, self-paced, un-cued, and freely moving, allowing us to use linear mixed effects models to examine measurements within the context of time and other emitted behaviors. We find that both recent lever press durations and a longer history of pressing contribute to decision-making. Importantly, oft-neglected variables such as checking for reward and the passage of time between actions have large contributions to decision-making (i.e. it is not just reward that drives decision-making). Thus, we were able to capture how one's subjective experience across a continuous space contributes to decision-making. This then allowed us to examine how corticostriatal circuits may represent and control this experience. Using in vivo calcium fiber photometry recordings of secondary motor cortex (M2) and its striatal projections, we found representation of subjective experience. Cell-type specific lesions and targeted optogenetic inhibition showed causal contribution to recent subjective experience and uncovered a distinct role for M2 striatal projections in the planning of upcoming actions based on experience. Thus, we identify behavioral mechanisms and neural correlates of subjective experience used to guide adaptive decision-making.

2-033. Spontaneous and sensory-evoked neuronal ensembles are distinct in awake mouse auditory cortex

Anton Filipchuk[1,2] ANTON.FILIPCHUK@CNRS.FR
Alain Destexhe[1] ALAIN.DESTEXHE@CNRS.FR
Brice Bathellier[3] BRICE.BATHELLIER@CNRS.FR

[1]Paris-Saclay Institute of Neuroscience (NeuroPSI)
[2]Integrative & Computational Neuroscience (ICN)
[3]Hearing Institute, Institut Pasteur

The brain is characterized by a prominent ongoing activity occurring in the form of large coactive groups of neurons defining neuronal ensembles which has been proposed as a general substrate for a wide variety of physiological computations and behaviors. For example, it was suggested, mostly based on local recording studies performed under anesthesia, that auditory cortex generates population patterns of spontaneous activity that resemble those observed under sensory stimulation, opening the possibility that cortical circuits may build a response to sensory stimuli using a pre-existing repertoire of assemblies expressed in spontaneous activity. To determine whether large scale sensory response in the awake auditory cortex is determined by pre-existent spatial neuronal patterns, we took advantage of large scale chronic two-photon calcium imaging techniques (field of view of 1x1 mm with up to 1200 neurons recorded simultaneously). We identified neural ensembles, including tens of neurons, activated both spontaneously and as a response to a large set of sounds. Comparing these population activity patterns using correlation metric we systematically observed similarities that were below the reproducibility of sound-evoked and spontaneous patterns. Thus spontaneous and evoked ensemble were largely distinct. This distinction was provided by small non overlapping fractions of neurons that were either predominantly spontaneously active or mostly responsive to the sounds, while the rest 2/3 of neurons were nonspecific. In contrast, during application of isoflurane anesthesia spontaneous and evoked ensembles became more stereotyped and were often very similar to each other. This convergence did not happen, however, for the thalamo-cortical input. Hence, dissociation of spontaneous and evoked activity patterns as seen in the awake cortex, is strongly sensitive to anesthesia, suggesting that the cortical circuit is able to produce different dynamical states with strikingly different coupling principles between internal dynamics and incoming information.

2-034. The causal role of the striatum in the encoding of task-adaptive expectation-based choice biases

Carlos Sindreu[1,2] CHASINDREU@YAHOO.COM
Rafael Marin[1] RAMARIN@CLINIC.CAT
Yerko Fuentealba[1] FUENTEALBA@CLINIC.CAT
Daniel Duque[3] DDDBCN@GMAIL.COM
Lorena Jimenez[1] LOJIMENEZ@CLINIC.CAT
Jaime de la Rocha[1] JROCHAV@CLINIC.CAT

[1]IDIBAPS
[2]cortical circuit dynamics
[3]Institut d'Investigacions Biomediques August Pi i Sunyer (IDIBAPS)

Learned expectations can influence perceptual decisions, but the underlying neural mechanisms have yet to be specified. We used a two-alternative auditory discrimination task that promoted the use of expectations by introducing serial correlation in the stimulus sequence in the form of repeating and alternating trial blocks (Fig. 1A). Recent work showed that rats leverage on these correlations to estimate the probability that the previous rewarded side is repeated and bias their choices accordingly, particularly when stimulus evidence is equivocal. A generative model that summed up the stimulus evidence, and the outcome-dependent running estimate of the first- and second-degree statistics of the sequence was able to accurately predict such flexible behavior (Fig. 1B). One brain region that may support the repeating bias is the associative striatum (i.e. DMS), as it encodes the value of the different options. Thus, we virally expressed the chloride-conducting opsin stGtACR2 in striatal projection neurons (Fig. 1C). Unilateral illumination during the whole trial generated a strong ipsilateral bias, confirming the role of the DMS in choice selection and the effectiveness of the method (Figs. 1D, E). We then hypothesized that silencing the DMS during the inter-trial interval, i.e. just before stimulus onset (Fig. 1F), would ablate the repetition bias. We obtained a real-time trial-by-trial estimate of the repeating bias and used it to perform closed-loop DMS inhibition while rats performed the task. Bilateral photostimulation occurred in half the trials with a high bias (~15% of total). Light-on trials showed reduced repetition bias compared with light-off trials, as evidenced by a marked decrease in the separation of the psychometric curves obtained in each of the blocks (Figs. 1G, H); no such effect was observed in controls (n=3 ACR2, n=2 GFP). We conclude that the dorsal striatum is important to maintain expectation priors that inform perceptual decisions.

2-035. Charting and navigating the space of solutions for recurrent neural networks

Elia Turner[1,2] ELIATURNER11@GMAIL.COM
Omri Barak[3] OMRI.BARAK@GMAIL.COM

[1]Technion - Israel Institute of Technology
[2]Computer Sciennce
[3]Technion, Haifa

In recent years, trained recurrent neural networks (RNNs) were used as models for neural activity of behaving animals. In this framework, networks are trained to perform a task that is similar to the experimental one. Features of the network's behavior and neural activity are then compared to neural recordings. Despite this increasing use, it is still unknown in which cases this approach will work. In particular, the match between model and data seems to suggest a unique solution found by both biology and the artificial network - a puzzling conjecture. %are a class of Machine learning models used to solve tasks with sequential structure. Since most neural circuits are recurrent, RNNs are often deployed to explain neural activity during computational tasks with a temporal dimension. Unlike their natural counterparts, an RNN's entire computation is accessible. Yet, most works treat them as a black box and analyze their generated activity directly. %A recent line of work uses RNNs as a hypothesis generation tool; The activity of a trained RNN is engineered backwards to reveal a low dimensional mechanism that may help explain the recorded activity. However, we still lack understanding of the hypothesis space itself. Recent work (Maheswaranathan et al., 2019) addressed the uniqueness problem, and proposed that the solutions to various canonical tasks are , from a topological perspective, widely universal. Here, we study a slightly more complex task – the Ready-Set-Go timing task. We find that universality no longer holds for this task, and that even identical settings can lead to qualitatively different solutions - both from behavioral and neuronal perspectives. We discover these differences by testing the trained networks ability to extrapolate, as perturbation to a system often reveals hidden structure. With the goal of understanding the solution space, we cluster the solutions into discrete sets and characterize each - showing that the neural and behavioral clusters are highly consistent. We draw a low-dimensional map of the solution space and sketch the training process as a trajectory in it. We suggest that the effects of updating parameters during training often takes form as a bifurcation in the governing discrete ODE, which results in a topological change in the dynamical mechanism and in the behavior. Moreover, we explore the question of nature vs. nurture - the effect of the initial weights vs. training set over the final solution and show that, in our setting, only the former has a meaningful impact over the learned solution.

2-036. Bandwidth expansion in the brain: Optimal encoding manifolds for population coding

Movitz Lenninger[1,2] MOVITZLE@KTH.SE
Mikael Skoglund[1] SKOGLUND@KTH.SE
Pawel Herman[1] PAHERMAN@KTH.SE
Arvind Kumar[3] ARVKUMAR@KTH.SE

[1]KTH - Royal Institute of Technology
[2]Intelligent Systems
[3]KTH Royal Inst of Technology

Stimuli in the brain are represented in the population activity of neurons, where individual neurons are tuned to respond to a small set of stimulus values. At the population level, tuning of neurons implies that every stimulus is mapped onto a D-dimensional encoding surface (D = dimensionality of the stimulus) embedded in an N-dimensional space (N= # of active neurons). Typically, N>>D. Mathematically, the representation of stimuli in the neural population activity is identical to the framework of bandwidth expansion in communication theory. Here, we exploit the bandwidth expansion framework to address the question: given noise and correlations, what is the optimal shape of the encoding surface and tuning curves? The notion of encoding surface led us to distinguish between two types of neural noise: channel noise (noise in spiking mechanism, background input) and observation uncertainty (noise in the stimulus-related input). Both noise sources can result in local (weak distortion) or global estimation errors (threshold distortion). We show that, in the case of channel noise, minimizing weak or threshold distortion leads to contradictory surface shapes. The optimal shape of the encoding surface is a trade-off between locally curved and globally flat. On the other hand, to minimize the effects of observation uncertainty, the optimal encoding surface should be flat. Thus, we propose that the optimal shapes of tuning curves depend on the relative strength of channel noise and observation uncertainty. We show that sparse coding (small tuning width) is suboptimal when observation uncertainty is large. These considerations also suggest that multi-peak tuning curves (grid cells) are more susceptible to threshold distortion than single peaked ones. Finally, taking the example of typical tuning curves seen in the early visual system, we show that threshold distortion becomes an acute problem when neurons have multi-modal tuning curves.

2-037. Dynamical regimes of electrically coupled inhibitory networks in the cerebellar cortex

Harsha Gurnani[1,2] HARSHA.GURNANI.16@UCL.AC.UK
N Alex Cayco Gajic[3,4] NATASHA.CAYCO.GAJIC@ENS.FR
R. Angus Silver[1] A.SILVER@UCL.AC.UK

[1]University College London
[2]Department of Neuroscience, Physiology and Pharmacology
[3]Ecole Normale Superieure
[4]Group for Neural Theory

Golgi cells (GoCs), the main inhibitory neurons in the cerebellar input layer, are thought to control the gain, threshold, and spike timing of vast numbers of downstream granule cells (GCs) through feedback and feedforward inhibition. Computational studies typically model GoC network activity as synchronous oscillations [Maex & De Schutter 1998] or chaotic dynamics [Rossert et al 2015], but despite extensive characterisation at the cellular level, there is little experimental evidence of how GoC network activity is organised. To address this, we used random-access acousto-optic-lens 3D two-photon microscopy to image calcium in local GoC populations in awake mice. Our results show that GoC population activity is organised into multiple modes during spontaneous behaviours. High pairwise activity correlations arising from widespread slow modulation of GoC network – a 'common mode' – reflect behavioural engagement. However, GoCs also exhibited heterogeneous activity – 'differential modes' – on faster timescales, that encode multidimensional behavioural activity. We observed similar multidimensional dynamics in experimentally constrained biophysical models of GoC networks, with a 'common mode' that was robust across different activity levels, clustering and dimensionality of their excitatory inputs, but critically depended on electrical coupling between GoCs. We further established that spatial dependence of coupling lies at a transition, with synchronous activity that undergoes transient desynchronization in response to external inputs [see also Vervaeke et al 2010]. Our results suggest that electrical coupling and its spatial dependence are critical determinants of GoC network dynamics, controlling the emergence of both network-wide and temporally

dispersed activity. This suggests that GoC- GrC inhibition is organised on distinct spatiotemporal scales, with global inhibition delivering adaptive gain control of GCs [Billings et al 2014, Cayco-Gajic et al 2017], and temporally precise and heterogeneous inhibition providing temporal patterning s [Litwin-Kumar et al 2017, Rossert et al 2015], potentially increasing the computational and learning capacity of the cerebellar cortex.

2-038. A synapse-centric account of the free energy principle predicts triplet STDP and stochastic synaptic release

David Kappel DAVID.KAPPEL@UNI-GOETTINGEN.DE
Christian Tetzlaff TETZLAFF@PHYS.UNI-GOETTINGEN.DE
Florentin Worgotter WORGOTT@PHYS.UNI-GOETTINGEN.DE

University of Goettingen

The free energy principle (FEP) is a mathematical formulation that describes how biological systems self-organize and survive in their environment. This principle has proven to work on multiple scales, from high-level behavioral and cognitive functions such as attention or foraging, down to the dynamics of specialized cortical microcircuits, suggesting that the FEP manifests on several levels of brain function. But despite its success, the FEP is often criticized for being too general making it hard or impossible to infer experimental predictions. Here, we apply the FEP to one of the smallest functional units of the brain: single excitatory synaptic connections. By constraining ourselves to a very limited and experimentally well understood biological system we are able to derive learning rules from first principles while keeping assumptions minimal. This synapse-centric account of the FEP predicts that synapses interact with the soma of the post-synaptic neuron through stochastic synaptic releases to probe their behavior and use feedback to update the synaptic weights. The emergent learning rules are regulated triplet STDP rules that depend only on the timing of the pre- and post-synaptic spikes and the internal states of the synapse. The parameters of the learning rules are fully determined by the parameters of the post-synaptic neuron model, suggesting a close interplay between the synaptic and somatic compartment and making precise predictions about the synaptic dynamics. We show that the FEP learning rules can be applied to large neural networks with thousands of synapses to solve a pattern classification tasks. This demonstrates that our approach can be scaled up to practical problems and the inherent local computation makes our learning rules also interesting for brain-inspired hardware.

2-039. Randomly mixed modular grid-place cell network for high capacity associative memory

Sugandha Sharma[1,2] SUSHARMA@MIT.EDU
Sarthak Chandra[1] SARTHAKC@MIT.EDU
Rishidev Chaudhuri[3,4] RCHAUDHURI@UCDAVIS.EDU
Mirko Klukas[1,5] MIRKO.KLUKAS@GMAIL.COM
Ila R. Fiete[1] FIETE@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Brain and Cognitive Sciences
[3]University of California, Davis
[4]Center for Neuroscience
[5]Department of Brain and Cognitive Sciences

How does the hippocampus, vastly outnumbered by the number of potential patterns in cortical neurons, store memories of experiences encoded by cortex? It is assumed to act as an associative memory, but classical models such as the Hopfield network store and robustly recall only up to N arbitrary patterns in a network of size N. Recently proposed associative memory architectures enable robust reconstruction of an exponential number of patterns; however, these networks involve biologically implausible assumptions about synapses or the amount of information recalled per pattern is small because they are highly structured.

We construct an entorhinal-hippocampal (EC-HC) attractor network using the theme of modular input structures with random feedforward EC-HC projections and associatively learned return weights that exhibits exponentially many robust (large-basin) fixed points. Training the HC to EC weights over only a small subset of contiguous grid cell states results in the formation of a large number of additional HC fixed points with large basins that can be reconstructed robustly from unseen grid cell states. In other words, these untrained hippocampal patterns can be robustly reconstructed in the presence of large amounts of noise from grid inputs. Further, the addition of hetero-associative recurrent synapses within HC leads to learning of the grid state transition matrix and the potential for high-capacity sequence memory.

Because place cells construct conjunctions between sensory experience and abstract spatial coordinates, we add a sensory input stream to HC. The combination of structured and unstructured inputs enables storage and robust recollection of a large number of arbitrary sensory patterns from partial cues. Thus, randomly mixed modular networks may support high-capacity associative memory, essential for interference-free episodic memory and spatial learning across time.

2-040. Using evolutionary algorithms to explore the dynamic operation of the CA1 hippocampal microcircuit

Andrea Navas-Olive[1] ACNAVASOLIVE@GMAIL.COM
Teresa Jurado-Parras[2] MTERESA.JURADO.PARRAS@GMAIL.COM
Liset Menendez de la Prida[2] LMPRIDA@CAJAL.CSIC.ES

[1]Cajal Institute
[2]Instituto Cajal

The hippocampus is critical for episodic memory function, but the basic mechanisms are elusive. During exploration, the neuronal firing of hippocampal CA1 pyramidal cells is organized by theta oscillations, a major hippocampal rhythm of 4-12 Hz. Traditionally viewed as a single layered structure, recent evidence suggests an exquisite laminar organization across deep and superficial pyramidal sublayers at the transcriptional, morphological and microcircuit levels [1,2]. Understanding how the intrinsic properties of different cell types and their connectivity constrain the organization of neuronal firing during theta is critical to better crack the hippocampal code.

Here, we exploit evolutionary algorithms to address the effect of single-cell heterogeneity on CA1 microcircuit operation. First, we developed a biophysically realistic model of CA1 pyramidal cells using the Hodgkin-Huxley multi-compartment formalism in the Neuron+Python platform and the morphological database Neuromorpho.org. We adopted genetic algorithms to identify a set of passive, active and synaptic conductances resulting in non-fitted realistic electrophysiological neural behavior during theta rhythm, including bimodality on the basal preferred phases. Then, we simulated the deep/superficial connectome resulting from local circuit interneurons as well as intra-hippocampal and entorhinal inputs and explored how changes of these factors modulated firing output [3].

By combining results from all simulations in a logistic regression model, we derived several predictions on the effect of up/down-regulation of different factors in constraining single cell firing during theta. Finally, we evaluated how the dynamic interplay between the CA3 and ECIII inputs, as well as cell-type specific GABAergic connections, can shift theta phases. Using high-density recordings in head-fixed mice running in a wheel, we tested these predictions experimentally. We discuss our results in the context of mechanisms underlying theta sequence dynamics [4, 5].

2-041. Animal movement reveals task-related information at different timescales

Chaoqun Yin[1,2] CHAOQUNYIN@G.UCLA.EDU
Alex Kostiuk[1] AKOSTIUK@MEDNET.UCLA.EDU
Simon Musall[3] SIMON.MUSALL@GMAIL.COM
Xiaonan Richard Sun[4] SUN.XIAONAN@ALUMNI.PRINCETON.EDU
Steven Gluf[5] SGLUF@CSHL.EDU
Anne Churchland[1] ACHURCHLAND@MEDNET.UCLA.EDU

[1]University of California, Los Angeles
[2]Neurobiology
[3]Institute of Biological Information Processing
[4]Zucker School of Medicine at Hofstra/Northwell
[5]Cold Spring Harbor Laboratory

Developments in automated movement tracking open the door to understanding how movements on multiple timescales relate to diverse aspects of behavior and internal state. These tools are especially relevant to decision-making: recent studies have revealed the preponderance of spontaneous, uninstructed movements, but have left unanswered how these movements relate to single trial choice/outcome or to slow fluctuations in performance. We aimed to better understand this relationship. We used a video dataset from 7 head-fixed mice that were trained to lick the spout corresponding to the spatial location in which the auditory clicks were presented. We deployed DeepLabCut to track the movements of 26 points labelled on the mice. We report three results. First, the movement pattern of animals tracked by DeepLabCut could predict the animal's upcoming choice, expected outcome and trial history. This indicates that videos contain rich task-related information which may be hardly perceptible to human observers. Second, we found that movements could discriminate the presence or absence of transiently presented optogenetic stimulation in parietal cortex. This was true although traditional movement metrics, like reaction time, were unaffected. Finally, we considered the variance remaining in the tracking after task-related information was regressed out. This quantity, termed, “task-independent variance” reflects movements that are unrelated to task variables we tested, and may reflect uninstructed fidgets or spontaneous exploration. Interestingly, in 6/7 animals, high task-independent variance was associated with lower overall performance and vice versa, indicating that spontaneous movements may reflect disengagement or may interfere with the animal's ability to do the task effectively. Taken together, our results highlight the insights offered to decision-making studies by movement tracking. Further, they reveal the importance of understanding slower-timescale fluctuations in movements that may reflect other processes and predict overall performance.

2-042. CRF neurons in a paraventricular thalamic circuit regulate food-approach vs. threat-avoidance conflict

Xu Zhang[1,2] XU.ZHANG.1@UTH.TMC.EDU
Douglas Engelke[1] DOUGLAS.S.ENGELKE@UTH.TMC.EDU
John O'Malley[1] JOHN.OMALLEY@UTH.TMC.EDU
Jose Fernandez-Leon[1] JOSE.A.FERNANDEZLEON@UTH.TMC.EDU
Sa Li[3] SA.LI@UMANITOBA.CA
Gilbert Kirouac[3] GILBERT.KIROUAC@UMANITOBA.CA
Michael Beierlein[1] MICHAEL.BEIERLEIN@UTH.TMC.EDU
Fabricio Do Monte[1] FABRICIO.H.DOMONTE@UTH.TMC.EDU

[1]University of Texas Health Science Center at Houston
[2]Neurobiology and Anatomy
[3]University of Manitoba

Balancing food-seeking with threat-avoidance behaviors is crucial for animals to survive, but the neural circuits that regulate this motivational conflict remain largely unknown. To answer this question, we designed an ethologically relevant “approach-food vs. avoid-predator threat” conflict test in which rats need to overcome their fear of predator odor to reach food. During the test, cat saliva was positioned in the food area where rats were trained to press a lever for sucrose in the presence of an audiovisual cue. Rats exhibited robust defensive behaviors and a clear suppression in food-seeking responses in the presence of the predator odor. Using in situ hybridization, immunohistochemistry and in vivo single-unit recordings from photoidentified cell-types, we identified a subpopulation of neurons in the anterior portion of the paraventricular thalamic nucleus (aPVT) which express corticotrophin-releasing factor (CRF) and are preferentially recruited during conflict. Chemogenetic inactivation of aPVTCRF neurons during conflict restored food-seeking behavior and reduced defensive responses, but had no effect when the predator odor or the food-seeking tasks were carried out independently. Using both anterograde and retrograde viral tracing methods, we characterized the anatomical connectivity between aPVTCRF neurons and brain regions that are implicated in the regulation of food seeking and defensive responses. aPVTCRF neurons project densely to the nucleus accumbens (NAc), and optogenetic activation of the aPVTCRF-NAc pathway recapitulated the predator-odor induced food-seeking suppression and avoidance responses by mediating target-dependent synaptic transmission in the NAc. In addition, we identified the ventromedial hypothalamus (VMH) as a critical input to aPVTCRF neurons, and demonstrated that aPVT-projecting VMH neurons are activated by predator odor and synapse onto NAc-projecting aPVT neurons. Chemogenetic inactivation of aPVT-projecting VMH neurons reduced defensive behaviors exclusively during conflict. Together, our findings describe a subpopulation of neurons in a hypothalamic-thalamostriatal circuit that suppresses reward-seeking behavior under the competing demands of avoiding threats.

2-043. Automatic tracking of mouse social posture dynamics by 3D videography, deep learning and GPU-accelerated robust optimization

Christian Ebbesen[1,2] CHRISTIAN.EBBESEN@NYUMC.ORG
Robert Froemke[1] ROBERT.FROEMKE@MED.NYU.EDU

[1]New York University
[2]Neuroscience Institute

Social interactions powerfully impact both the brain and the body, but high-resolution descriptions of these important physical interactions are lacking. Currently, most studies of social behavior rely on labor-intensive methods such as manual annotation of individual video frames. These methods are susceptible to experimenter bias and have limited throughput. To understand the neural circuits underlying social behavior, scalable and objective tracking methods are needed. We present a hardware/software system that combines 3D videography, deep learning,

physical modeling and GPU-accelerated robust optimization. Our system is capable of fully automatic multi-animal tracking during naturalistic social interactions and allows for simultaneous electrophysiological recordings. We capture the posture dynamics of multiple unmarked mice with high spatial (~2 mm) and temporal precision (60 frames/s). Our method is based on inexpensive consumer cameras and is implemented in python, making our method cheap and straightforward to adopt and customize for studies of neurobiology and animal behavior. We are currently combining our method with recordings of single neurons in the somatosensory cortex of freely interacting mice. Using our 3D tracking method, we can automatically relate the firing patterns of these neurons to social events (nose-to-nose touch, anogenital sniffing, mounting), posture and movement kinematics, postures and movement displayed by the social interaction partner, and spatial features (heading direction, spatial location).

2-044. Average reward rate shapes both tonic and phasic dopamine activity during patch foraging

Laura Grima[1] GRIMAL@JANELIA.HHMI.ORG
Thomas Akam[2] THOMAS.AKAM@PSY.OX.AC.UK
Mark Walton[2] MARK.WALTON@PSY.OX.AC.UK

[1]HHMI Janelia Research Campus
[2]University of Oxford

Decision neuroscience has typically focused on how organisms choose amongst simultaneously presented options. However, foraging animals typically encounter options sequentially, and must decide whether to stay with the current option or search for an alternative. This makes the opportunity cost of time critical; the animal must weigh the expected rate of return from the current option against the average rate across the environment. One popular theoretical account proposes that tonic dopamine activity encodes the average reward rate. However, there have been few direct tests of this hypothesis, nor is it known whether and how average reward rate affects phasic responses to reward predicting cues. We addressed this question by training mice in a patch foraging task where they chose between harvesting rewards from the current depleting patch or travelling to a new patch. The richness of individual patches varied pseudorandomly from patch to patch, while the environmental reward rate was independently manipulated by varying the travel time between patches. Patch leaving decisions were qualitatively consistent with optimal foraging theory, with a positive influence of the current patch's richness, and negative influence of environment richness, on number of rewards harvested before leaving. We recorded calcium activity in VTA dopamine neuron cell bodies, and axon projections in the nucleus accumbens core (NAcC), using fiber photometry (n=14 mice, ~10,000 patches). Baseline, putatively tonic, dopamine activity was positively modulated by both local reward rate in the current patch and average reward rate in the environment. However, phasic activity was also independently shaped by the environment; shorter travel times – higher environment reward rates – potentiated reward cue responses. These data support the hypothesis that average reward rate shapes tonic dopamine, but reveal a previously unknown modulation of phasic cue responses by average environment reward rate, placing dopamine in a position to guide 'stay or leave' decisions.

2-045. Long-term turnover dynamics in area CA1 of hippocampus are consistent with plasticity of non-spatial inputs

Federico Devalle[1] FEDERICO.DEVALLE@GMAIL.COM
Alon Rubin[2] ALON.RUBIN@WEIZMANN.AC.IL
Yaniv Ziv[2] YANIV.ZIV@WEIZMANN.AC.IL
Alex Roxin[1] AROXIN@CRM.CAT

[1]Centre de Recerca Matematica
[2]Weizmann Institute of Science

Recent chronic imaging experiments in mice have revealed that the hippocampal code exhibits non-trivial turnover dynamics over long time scales [1]. Specifically, the subset of cells which are active on any given session in a familiar environment changes over the course of days and weeks. While some cells transition into or out of the code after a few sessions, others are stable over the entire experiment. The mechanisms underlying this turnover are unknown. Here we show that the statistics of turnover are consistent with a model in which non-spatial inputs to CA1 pyramidal cells readily undergo plasticity, while spatially tuned inputs are largely stable over time. The heterogeneity in stability across the cell assembly, as well as the decrease in correlation of the population vector of activity over time, are both quantitatively fit by a simple toy model with Gaussian input statistics. In fact, such input statistics emerge naturally in a network of spiking neurons operating in the fluctuation-driven regime. This correspondence allowed us to map the parameters of a large-scale spiking network model of CA1 onto the simple statistical model, and thereby fit the experimental data quantitatively. Our model suggests that the internal representation of space in the hippocampus evolves over time due to changes in non-spatial inputs, which may represent changing contextual cues, or simply the passing of time. It also suggests that the locus of plasticity underlying turnover may be in the inputs from the entorhinal cortex, and not necessarily CA3.

2-046. Frontal eye field and caudate neurons make different contributions to reward-biased perceptual decisions

Yunshu Fan[1] FANYUNSHU@GMAIL.COM
Joshua Gold[2] JIGOLD@PENNMEDICINE.UPENN.EDU
Long Ding[2,3] LDING@PENNMEDICINE.UPENN.EDU

[1]Brown University
[2]University of Pennsylvania
[3]Department of Neuroscience

Complex decisions often require interpreting external sensory inputs in the context of outcome expectations. Previous studies suggested that the caudate nucleus of the basal ganglia and one of its major cortical inputs, the frontal eye field (FEF) of the lateral prefrontal cortex, may contribute to saccade decisions that balance visual evidence and reward expectation. To better understand these contributions, we compared single-neuron activity in the caudate nucleus and FEF in three monkeys performing a two-alternative reaction-time random-dot visual motion discrimination task. Across trials, we manipulated visual evidence by presenting motion stimuli with varying strengths. Across blocks of trials, we manipulated reward expectation by assigning large reward for one choice and small reward for the other. We also compared neural activities in the two regions to accumulate-to-bound model predictions derived from behavior. We found that both regions share coarse similarities: (1) a substantial fraction of neurons in both regions were sensitive to both stimulus properties and the reward-choice association; (2) the average activity profile of the choice-sensitive neurons in both regions were consistent with evidence accumulation. However, we also identified inter-regional differences, including quantitative differences in the prevalence of neurons selective for various task factors and in how the neural activity directly relates to predictions from fits of a drift-diffusion model (DDM) to behavior. Specifically, certain DDM predictions, including a "bound-crossing"-like pattern and a reward context-dependent bias in drift rates, were evident in FEF but not caudate. Previously, a study [1] found that the micro-stimulation effects of caudate could be better explained by modulating the coordination between drift-rate bias and bound bias in DDM than affecting individual DDM components directly. This, together with our results, suggest that FEF is more directly involved than caudate in implementing certain components of a DDM-like decision, whereas caudate's roles are likely modulatory in nature.

2-047. Adult neurogenesis improves spatial information encoding in the hippocampus

M. Agustina Frechou[1,2] MFRECHOU@MAIL.EINSTEIN.YU.EDU
Sunaina Martin[1] SUNAINA.MARTIN@EINSTEIN.YU.EDU
Ruben Coen-Cagli[3] RUBEN.COEN-CAGLI@EINSTEINMED.ORG
J. Tiago Goncalves[1] TIAGO.GONCALVES@EINSTEINMED.ORG

[1]Albert Einstein College of Medicine
[2]Neuroscience
[3]Albert Einstein College Of Medicine

The dentate gyrus (DG) is one of the few regions of the mammalian brain that continue to add new neurons through adulthood. Adult-born neurons (ABNs) develop in the DG into mature granule cells which are physiologically indistinguishable from developmentally born neurons. ABNs are active and contribute to memory discrimination tasks during a critical developmental period at 4-6 weeks post-mitosis. Exposing animals to enriched environments (EE), consisting of large cages with toys, domes and running wheels, results in an increase in the number of ABNs, as well as improvement in context discrimination and spatial memory. Conversely, decreasing the number of ABNs results in spatial memory deficits. However, it is unknown how a small cohort of ABNs can influence the network activity and information encoding of the DG. Here, we demonstrate that exposing mice to EE results in an increase in the spatial information encoding of the DG neuronal population responses compared to controls, and that this increase can be largely attributed to sharper spatial tuning of individual DG neurons. Conversely, ablating DG adult neurogenesis resulted in a decrease in spatial tuning. Importantly, the increase in spatial tuning seen in EE mice is not present in mice exposed to EE where neurogenesis was ablated by focal irradiation, indicating that ABNs are necessary for the effects of EE on DG network activity. Additionally, we show that acutely silencing ABNs using a chemogenetic approach, results in EE-exposed animals being less able to discriminate between similar contexts during a contextual fear conditioning task, indicating that ABN activity is

required, at least partially, for the task. These results demonstrate, for the first time, that ABNs increase both the information content of the DG and the spatial tuning of individual DG granule cells.

2-048. Neural mechanism for predicting the auditory consequences of behavior

Helen Hou[1] HELENXUNHOU@GMAIL.COM
Naz Taskin[1] NAZTASKIN@GMAIL.COM
Caroline Yu[1] CY2392@COLUMBIA.EDU
Michael Genecin[1] MG4178@COLUMBIA.EDU
Ioannis Kymissis[1] JOHNKYM@EE.COLUMBIA.EDU
Larry Abbott[1] LFA2103@COLUMBIA.EDU
Nathaniel B Sawtell[2] NS2635@COLUMBIA.EDU

[1]Columbia University
[2]Zuckerman Mind Brain Behavior Institute, Columbia University

The brain-body interaction is a two-sided coin: the brain controls the body to generate behaviors, but behavior itself can pose challenges for sensory function. Sounds generated by our own movements, such as eating crunchy food, can produce sensory responses that interfere with the detection of behaviorally relevant sounds, such as the footsteps of a predator. Past work has shown that the first central stage of auditory processing in mammals, the dorsal cochlear nucleus (DCN), selectively cancels responses to self-generated sounds1. Here we combine experiments and modeling to explore the circuit mechanisms underlying this cancellation. We developed a preparation that allows for high-resolution tracking of facial kinematics and stable single-unit recordings in the brainstem during eating in head-fixed mice. We discovered that a major class of inhibitory interneuron in the DCN, known as cartwheel cells (CWCs), exhibit non-auditory responses with temporal profiles that match those of eating-related auditory input and hence could contribute to cancelling self-generated sounds in the output cells of the DCN to which they project. CWCs share numerous similarities with Purkinje cells of the cerebellum and Purkinje-like medium ganglion cells of the electrosensory lobe, including similar forms of synaptic plasticity at excitatory parallel fiber synapses. Unlike Purkinje and medium ganglion cells, however, CWCs lack an instructive signal for guiding plasticity at parallel fiber synapses2. We develop a model constrained by data from in vitro and in vivo recordings illustrating how cancellation signals could be constructed in CWCs in the absence of a teaching signal.

2-049. Dynamics of drifting receptive fields during noisy representation learning

Shanshan Qin SSQIN@SEAS.HARVARD.EDU
Cengiz Pehlevan CPEHLEVAN@SEAS.HARVARD.EDU

Harvard University

Long-term memories and learned behavior are conventionally associated with stable neuronal representations. However, recent experiments showed that neural population codes in many brain areas continuously change even when animals have learned and stably perform their tasks. This representational “drift” naturally leads to questions about its causes, dynamics, and functions. To address these questions, we consider a basic and biologically-plausible model of representation learning, the Hebbian/anti-Hebbian network, which learns localized receptive fields (RFs) that tile the input data manifold. We hypothesize that drifting representation is due to noisy synaptic updates and turnover. We find that while the receptive fields of individual neurons change significantly over time, the representational similarity of population codes is stable over time, providing a substrate for potential robust downstream readout. We investigate the dynamics of the drift and find that the relative stability of RFs is correlated with the strength of tuning. Our model accounts for many of the recent experimental observations in the hippocampus and cortices related to sensorimotor tasks.

2-050. Dynamic encoding of saccade sequences in primate frontal eye field

Jing Jia[1,2] cherry930908@163.com
Zhen Puyang[1] wxpzz0365@163.com
Qingjun Wang[1] 2573396945@qq.com
Xin Jin[3] xjin@salk.edu
Aihua Chen[1] ahchen@brain.ecnu.edu.cn

[1]East China Normal University
[2]Key Laboratory of Brain Functional Genomics
[3]Salk Institute for Biological Studies

Frontal eye field (FEF) is a key part of oculomotor system, and it is known as a critical area for saccade movement in primate (Schall JD 2002). However, whether and how FEF contributes to sequential saccades remain largely unknown. To investigate these questions, we trained rhesus macaques to perform a novel self-paced oculomotor sequence task consisting of four (Left-Left-Right-Right) consecutive saccades, which allow a good internal control for studying the sequence vs. direction-related activities (Geddes CE et al. 2018). By recording the neuronal activities in FEF of two monkeys (226 well isolated neurons), we found that FEF uses a dynamically encoding strategy for sequential saccade task, and some FEF neurons show activities related with the initiation, switch and termination of the sequence. In addition, the sequence-related activities are context-dependent, with different firing activities during memory- versus visually-guided sequence. To further investigate the dynamic coding characteristics and the causal role of FEF during sequential saccades, electric microstimulation and reversible inactivation in FEF during the performance of sequential saccades were applied, and the results turned out that the supra-threshold microstimulation in FEF evokes saccade but does not affect the overall sequence structure, and interestingly the same direction of saccade at different positions within the sequence was affected differently, indicating that the effect of microstimulation-triggered saccades are more complicated than just generating a saccade to certain direction. Reversible inactivation of FEF severely impaired the performance of sequence, also with the effect varies for the different action within sequence, rather than only affecting the saccade movement along certain direction. These results reveal the context-dependent dynamic encoding of saccade direction in FEF, and underscore a critical role of FEF in planning and execution of sequential saccades.

2-051. Brain map: A multiscale connectome derived from whole-brain volumetric reconstructions

Christopher Brittin[1] cabrittin@protonmail.com
Steven Cook[2] sjc2153@columbia.edu
David Hall[3] david.hall@einstein.yu.edu
Scott Emmons[3] scott.emmons@einstein.yu.edu
Netta Cohen[4] n.cohen@leeds.ac.uk

[1]Memorial Sloan Kettering Cancer Center
[2]Columbia University
[3]Albert Einstein College of Medicine
[4]University of Leeds

A primary goal of systems neuroscience is to understand how the brain's structure and function combine to generate behavior. Achieving this goal is complicated by the fact that structure and function are intertwined across scales; from the nano-level localization of synapses, through multiplicities on neuronal morphologies and their contribution to circuit organization, to the high-level stereotyped connection between different regions of the brain. Such nested complexity means that to date, we have yet to reconstruct and model the structure of a complete nervous system that is integrated across all of these scales. Here, we present a complete structure-function model of the nematode C. elegans main neuropil, the nerve ring, which we derive by integrating the volumetric reconstruction from two animals with corresponding synaptic and gap junctional connectomes. By spatially grounding our analysis within the volumetric data, we are able to demonstrate that the C. Elegans connectome is not invariant, but that a precisely wired circuit is embedded within a background of variable connectivity, and propose a corresponding reference connectome for the core circuit. Whereas previously the nerve ring was considered a densely packed tract of neurites, we are able to use the reference connectome to demonstrate that the high-level spatial organization suggests modular circuits, with distinct functional roles. Modules are comprised of predominantly local and some cross-modular neurons with distinct micro- and nanoscale features that support brain-wide coordination. By combining these features in the data, we are able to posit a 3-layered brain map of the nerve ring, where parallel information processing modules converge onto a recurrent, highly distributed sub-circuit. The brain map is reminiscent of the layered connectivity of pyramidal neurons in the mammalian cortex and its biologically inspired analog – Residual Networks – pointing to robust features of brain organization that likely scale to larger nervous systems.

2-052. Behavioral modulation of information processing in visual cortex analysed with Copula-GP model

Nina Kudryashova[1,2] NKUDRYAS@ED.AC.UK
Theoklitos Amvrosiadis[1] T.AMVROSIADIS@ED.AC.UK
Nathalie Dupuy[1] NATHALIE.DUPUY@GMAIL.COM
Nathalie Rochefort[1] N.ROCHEFORT@ED.AC.UK
Arno Onken[1] AONKEN@INF.ED.AC.UK

[1]University of Edinburgh
[2]Informatics

In previous work, visual information processing in the primary visual cortex was shown to be strongly modulated by behavior (e.g. running speed) and environmental context (e.g. different variants of the task). Analysis of this modulation is crucial for understanding sensory integration and adaptation to the changes in the environment. Yet, this analysis proves difficult due to high dimensionality of the data, as well as different latent timescales (spanning milliseconds to days) and different statistics of the recorded variables (e.g. discrete or continuous). We have developed a method for statistical multivariate analysis of neuronal and behavioral data that combines scalability to high dimensions, accuracy in entropy estimation, and interpretability of the model components. The method is based on parametric copulas which separate the statistics of the individual variables from their dependence structure, allowing for accurate mutual information estimation. We use Gaussian Process (GP) priors over copula parameters, conditioned on a continuous task-related variable. This parameterization makes the contextual dependence explicit and accounts for the uncertainty in model parameters. Finally, we combine bivariate models in vine copula constructions, enabling the analysis of high dimensional datasets. After validation on synthetic data, we apply the method to two-photon calcium imaging data of neuronal population activity in the primary visual cortex of awake mice engaged in a visuospatial navigation task in a virtual reality environment. We show that our Copula-GP model captures significantly more information in the high-dimensional neuronal data compared to the other commonly used information estimators. Finally, we demonstrate that our model can eliminate the confounding information between the velocity and neuronal signals and reveal the navigational task structure in an unsupervised manner. We conclude that the Copula-GP framework enables comprehensive analysis of the information processing in large neuronal populations under behavioral and/or contextual modulation.

2-053. An image-computable ideal observer model of navigation explains homing behavior and grid/place field deformation

Yul Kang[1,2] YUL.KANG.ON@GMAIL.COM
Daniel M. Wolpert[3] WOLPERT@COLUMBIA.EDU
Mate Lengyel[1,2] M.LENGYEL@ENG.CAM.AC.UK

[1]University of Cambridge
[2]Department of Engineering
[3]Zuckerman Mind Brain Behavior Institute, Columbia University

Spatial navigation relies on vision to sense remote landmarks. However, vision only provides noisy and partial information and it is unknown how the resulting visual uncertainty affects navigational performance and neural representations. Here, we show that visual uncertainty underlies key effects of environmental geometry on spatial navigation and grid/place field deformations. We develop an image-computable ideal observer, which continually updates its probabilistic beliefs about its location by optimally combining visual and self-motion inputs via Bayesian filtering, thus making directly testable predictions for navigational behavior. The same ideal observer also predicts neural representations under Fisher-optimal population coding of location uncertainty. Critically, such codes exhibit potentially anisotropic expansion-contraction of grid/place fields depending on how location uncertainty changes along different directions as environmental geometry changes. Through analytical arguments we show how environmental geometry affects location uncertainty. We then use numerical simulations of the ideal observer following realistic trajectories under the conditions used in the relevant experiments. Our model correctly predicts that (1a) behaviorally, humans are less accurate when homing to the narrow side of a trapezoid-shaped room (compared to the broad side or a square-shaped room; Bellmund et al. 2020); (1b) neurally, rodent grid fields expand and lose symmetry on the narrow side of a trapezoid-shaped room (Krupic et al. 2015); (2a) behaviorally, when homing in a deformed room, humans' response distribution is stretched or compressed along the deformed axis and, paradoxically, shows the opposite pattern along the unchanged axis (Chen et al. 2015; Hartley et al. 2005); (2b) neurally, rodent grid/place fields show a similar pattern of deformation (O'Keefe & Burgess 1996; Barry et al. 2007). Thus, by connecting visual uncertainty, navigational uncertainty, and neural representations, our model provides a unifying principle underlying a broad range of phenomena about how homing behavior and grid/place fields depend on environmental geometry.

2-054. Mirror-symmetric viewpoint selectivity in deep neural networks

Amirhossein Farzmahdi[1,2] AFARZMAHDI@MAIL.ROCKEFELLER.EDU
Tal Golan[3] TAL.GOLAN@COLUMBIA.EDU
Reza Ebrahimpour[4] REBRAHIMPOUR@SRTTU.EDU
Winrich Freiwald[5] WFREIWALD@MAIL.ROCKEFELLER.EDU

[1]Institute for research in Fundamental Sciences
[2]School of cognitive Sciences
[3]Zuckerman Institute, Columbia University
[4]Shahid Rajaee University
[5]Rockefeller University

Face information is processed by a network of interconnected face-selective areas. Each face area is highly selective for faces, but represents face information in different ways. In the AL face patch, cells are tuned to mirror-symmetric viewpoints (e.g., left and right profiles) similarly. How is this invariance achieved? Here, we suggest a simple bottom-up explanation for mirror-symmetric viewpoint selectivity. First, we show that such mirror-symmetric viewpoint selectivity for faces emerges in the fully connected layer of convolutional deep neural networks trained on object classification such as AlexNet and VGG. Using a diverse set of 3D objects rendered from multiple views as stimuli, we show that such emergence of mirror-symmetric viewpoint selectivity is not unique to faces: it occurs for multiple other object categories with bilateral symmetry. To understand the causes of the emergence of this invariance, we dissect the response properties of the channels that contribute to it. We find that mirror-symmetric viewpoint selectivity emerges when reflection-invariant local features are spatially pooled by units with sufficiently large receptive fields. When a 3D object is bilaterally symmetric, such pooling operation leads to a representation that is identical for reflected viewpoints (e.g., left and right profiles). Moreover, the presence of bilateral symmetric object categories in the training set is critical for the formation of reflection-invariant features. Consequently, deep neural networks trained on datasets with mainly asymmetric object categories do not show mirror-symmetric viewpoint selectivity for faces or other object categories. We hypothesize that mirror-symmetric viewpoint selectivity emerges in biological brains for the same reasons it emerges in convolutional deep neural networks. If this hypothesis is correct, then mirror-symmetric viewpoint selectivity in the primate visual system should occur for multiple object categories. Such invariances can be uncovered by appropriate electrophysiological experiments.

2-055. Drifting assemblies for persistent memory

Felipe Yaroslav Kalle Kossio[1] YKALLEKO@UNI-BONN.DE
Sven Goedeke[1] SGOEDEKE@UNI-BONN.DE
Christian Klos[1] CKLOS@UNI-BONN.DE
Raoul-Martin Memmesheimer[1,2] RM.MEMMESHEIMER@UNI-BONN.DE

[1]University of Bonn
[2]Neural Network Dynamics and Computation, Institute for Genetics

The weights of biological neural networks and the neural representations they encode can change over time. Weight changes are activity dependent and thus assumed to be crucial for learning, but they also occur spontaneously, independent of previous activity and even in its absence. A similar dichotomy exists for neural representations: they change due to learning, but also spontaneously and even without altering behavior. In a standard model, assemblies of strongly interconnected neurons form the building blocks of associative memories. To faithfully store memories these assemblies are assumed to consist of the same neurons over time. Inspired by experimentally observed changes of connections and various neural representations, we propose that assemblies may change completely over time. We demonstrate this using different models, including networks of leaky integrate-and-fire and binary neurons. In our models, spontaneous turnover of synaptic connectivity or randomness of spiking induce exchanges of neurons between the assemblies. Activity dependent and homeostatic plasticity conserve the assembly structure, while the underlying neurons gradually change. The resulting “drift of assemblies” is compensated by similar unsupervised processes to keep the inputs, outputs and representations consistent, allowing persistent memories. Our model explains the results of recent experiments, which examined the evolution of a fear memory ensemble in the prelimbic cortex. Furthermore, we show that drifting assemblies enable non-trivial information processing. Our findings suggest that one needs to understand memory systems in their completeness, since their individual parts may strongly change, while memories persist due to unsupervised compensation of these changes. They also highlight probable mechanisms driving the changes of memory representations and enabling the stability of inputs and readouts.

2-056. Divisive normalization is an efficient encoder for log-logistic stimulus environments

Stefan Bucher WEB@STEFAN-BUCHER.CH
Adam Brandenburger ADAM.BRANDENBURGER@STERN.NYU.EDU

New York University

Divisive normalization (DN) is a canonical computation (Carandini&Heeger 2012) reputed to solve the brain's problem of representing potentially unbounded stimuli via bounded firing rates. Originally discovered in the primary visual cortex, DN has since been observed in many other sensory domains; it is often considered to be an implementation of the efficient coding principle, as it has been shown empirically to reduce redundancy in natural stimulus statistics (Schwartz&Simoncelli 2001). To date, however, it has never been analytically demonstrated that divisive normalization is provably efficient. We close this gap with a theoretical result that states the conditions under which divisive normalization is an optimal encoding system, and specify a class of input distributions that DN encodes efficiently. We show that in a low-noise environment DN maximizes the mutual information between an n-dimensional stimulus and its representation if and only if the distribution of the stimuli in the environment is multivariate log-logistic. Divisive normalization transforms (only) this particular distribution into an entropy-maximizing distribution, which amounts to a multivariate analog of histogram equalization. We extend this result to allow for an arbitrary metabolic cost of the representation, and show how this impacts the optimally encoded distributions. Our result suggests that divisive normalization may have adapted to efficiently representing such log-logistic stimulus distributions, which are indeed ecologically relevant: They bear a striking resemblance to the scale-invariant power law distributions occurring in many natural contexts ranging from the frequency spectra of natural images (Field 1987) to Benford's Law (according to which the leading digit of many naturally occurring numbers is likely to be small). Beyond formalizing the sense in which this important computation adheres to the efficient coding principle, our theoretical result also yields empirically testable predictions across sensory domains on how the DN parameters should be tuned to features of the input distribution.

2-057. VIOLA: Online spike inference for voltage imaging data analysis

Changjia Cai[1] CHANGJIA@LIVE.UNC.EDU
Cynthia Xiao Dong[1] XCYNTHIA@LIVE.UNC.EDU
Eftychios Pnevmatikakis[2] EPNEVMATIKAKIS@FLATIRONINSTITUTE.ORG
Marton Rozsa[3] ROZSAM@JANELIA.HHMI.ORG
Andrea Giovannucci[1,4] AGIOVANN@EMAIL.UNC.EDU

[1]University of North Carolina at Chapel Hill
[2]Center for Computational Mathematics Flatiron Institute
[3]HHMI Janelia Research Campus
[4]Joint Department of Biomedical Engineering

While optical methods, genetically encoded voltage indicators, and optogenetics already enable millisecond readout and control of large neuronal populations using light, the lack of corresponding advances in computational algorithms has slowed progress in closing the loop between brain and machine. The fundamental challenge is to reliably extract spikes (neural activity) from fluorescence imaging frames at speeds approaching brain information processing. The latest generation of voltage indicators have enabled readout of spikes with millisecond time resolution. However, no framework for the online extraction of spikes and subthreshold signals from voltage imaging dataset exists, and available calcium imaging algorithms are limited to ~30Hz maximum speed on standard datasets sizes. In order to tackle these challenges, we propose the first framework for streaming analysis of voltage imaging datasets, VIOLA, which exploits computational graphs and accelerated hardware to achieve unprecedented processing speeds. We provide optimized motion correction, source extraction and spike detection routines which achieve performance similar to batch algorithms while operating at hundreds of Hz in streaming. We compare VIOLA and VolPy - a batch algorithm for voltage imaging data analysis - on ground truth electrophysiology data, demonstrating comparable performance. VolPy's timing performance demonstrates that it can consistently sustain speeds of 280Hz on 256x256 datasets, about tenfold faster than existing streaming algorithms for calcium imaging data. Crucially, the frequent upgrades and developments in GPU hardware and TF will directly apply to our framework. Our work provides the first algorithm able to process streaming data from voltage imaging experiments, at speeds that are unprecedented and projected to rapidly approach real-time data collection (400-1000Hz). In conclusion, this work lays the foundation for interfacing large neuronal populations and machines using voltage imaging in real-time, enabling new applications in neuroprosthetics, brain-machine interfaces, and experimental neuroscience.

2-058. Transformation in feature representation under crowding between V1 and V4 neuronal populations

Christopher Henry[1,2] CHRISTOPHER.AARON.HENRY@GMAIL.COM
Adam Kohn[3] ADAM.KOHN@EINSTEINMED.ORG

[1]Albert Einstein College of Medicine
[2]Neuroscience
[3]Albert Einstein College Of Medicine

Visual features are encoded in a distributed manner spanning numerous cortical areas, with transformations in neuronal selectivity and cortical representation occurring along the visual hierarchy. To understand this distributed sensory code and how it underlies perception, we must identify how feature representations generalize across changes in the environment. Visual crowding, whereby judgments of target features are impaired by adjacent distractor stimuli, provides a powerful paradigm to address these questions, as it involves a range of spatial scales of integration and thus places profound constraints on the relationship between neuronal representations and perception.

We asked how crowding by distractor stimuli altered the representation of target orientation in simultaneously recorded neuronal populations in macaque V1 and V4. Using a population decoding approach, we found that crowding resulted in marked impairments in target orientation discriminability in both cortical areas. Information losses under crowding were more pronounced in V4 populations and varied over a larger range of spatial scales than in V1. Information loss occurred because crowding modulated neuronal responsivity and variability, with both response suppression and facilitation under crowding limiting feature encoding. While both cortical areas exhibited this diversity in modulation, V4 neurons were biased towards facilitation under crowding, in part due to their larger spatial receptive fields.

Among well-tuned V4 neurons, target tuning was altered more between crowded and uncrowded conditions than that in V1. Further, small changes in the orientation or location of distractor stimuli perturbed tuning more in V4, indicating greater configuration-dependence. Population-level analyses reinforced these failures of generalization in V4 and showed that target detection amid crowded displays was also impaired. These results reveal that crowding alters the V1 representation and that these effects are compounded in V4. Together these limits on feature discriminability closely approximate those seen in human perception as measured using similar displays.

2-059. Correlated variability of stimulus and action codes track internal models for perceptual choice

Ruud van den Brink[1,2] RUDYVDBRINK@GMAIL.COM
Keno Hagena[1] KENOHAGENA@GOOGLEMAIL.COM
Niklas Wilming[1] NIKLAS.WILMING@GMAIL.COM
Peter Muphy[3] MURPHYP7@TCD.IE
Tobias Donner[3] T.DONNER@UKE.DE

[1]University Medical Center Hamburg-Eppendorf
[2]Department of Neurophysiology
[3]murphyp7@tcd.ie

Many choice tasks entail arbitrary rules that govern the mapping from states of the sensory environment (e.g. orientation of a grating stimulus) to motor response (e.g., left- or right-hand movement). Animals and humans can flexibly change this sensory-motor mapping following rule changes. This key element of cognitive behavior may require the flexible configuration of networks of neurons encoding stimulus features and motor actions. Here, we show that the spontaneous, sustained, co-variation of neural population codes for sensory features and motor responses tracks changing internal representations of sensory-motor mapping rules. Human participants discriminated the orientation (horizontal vs. vertical) of peripheral visual gratings and reported their choices with left- or right-hand responses. The stimulus-response mapping rule alternated after variable numbers of trials. In different blocks of trials, the rule changes were instructed explicitly or had to be inferred from a stream of ambiguous cues around fixation. For the latter blocks, we fit a normative model to participants' behavior to reconstruct the time course of their belief about the active rule. We used fMRI multi-voxel pattern analysis to extract time courses of intrinsic fluctuations of orientation-selective activity in visual cortical regions and of action-selective activity in (pre-)motor cortical regions. Intrinsic fluctuations were measured between trials of the perceptual choice task. Stimulus- and action-selective activity fluctuations were correlated, with a sign that flipped along with the rule. When the rule had to be inferred under uncertainty, the magnitude of their correlation dynamically tracked the graded strength of the participants' belief about the active rule. Our findings show that dynamic changes in the co-fluctuation of sensory and motor population codes across cortex reflect the ongoing re-configuration of networks that is required for flexible decision-making. They also support the idea that internal task models are encoded in correlated neural variability.

2-060. Effect of dendritic nonlinearities on single-neuron pattern storage

Jacob Zavatone-Veth[1,2] JZAVATONEVETH@G.HARVARD.EDU
Cengiz Pehlevan[1] CPEHLEVAN@SEAS.HARVARD.EDU

[1]Harvard University
[2]Physics

Single-neuron information processing depends critically on richly nonlinear dendritic signal propagation. However, the computational role of dendritic nonlinearities is incompletely understood, as most theoretical analyses have focused on point-like neural models. As a tractable setting for the study of computation in branched neurons, Poirazi, Mel, and colleagues proposed a two-layer model in which each dendrite is itself a linear-nonlinear unit, the outputs of which are integrated at the soma. Though the effect of dendritic nonlinearities on the properties of such models has been studied in a few special cases, a detailed understanding of their role in determining expressive power has not emerged. Here, we comprehensively characterize how dendritic activation functions affect the pattern storage capacity of treelike two-layer models in which the dendrites have disjoint receptive fields. Using analytical methods from statistical physics, we show that the storage capacity is related to the smoothness of the dendritic activation functions. Non-smooth dendritic nonlinearities, which mimic dendritic spikes, yield non-bounded capacities in the limit in which the number of dendritic branches is large. In contrast, smooth activation functions, such as a sigmoid or a rectifier, yield finite capacities even in the many-branched limit. For such activation functions, our theory provides simple estimates of the pattern storage capacity. In general, we predict that nonlinearity should increase capacity, but may reduce the robustness of classification. Furthermore, we characterize how the storage capacity changes in the limit of very sparsely active input synapses. These connections between expressive power and smoothness begin to shed light on the influence of dendritic nonlinearity on single-neuron computation, and generate an intriguing hypothesis for the functional significance of dendritic spikes.

2-061. Low-dimensional shared variability disrupts communication between brain areas

Olivia Gozel GOZEL@UCHICAGO.EDU
Brent Doiron BDOIRON@UCHICAGO.EDU

University of Chicago

Neuronal activity in cortical brain areas is modulated by both stimulus structure as well as cognitive state, and ranges from asynchronous spiking to richly patterned spatio-temporal activity. Though activity across a neuronal population is varied and heterogeneous, group activity is often confined to a low-dimensional space. A growing number of studies characterize this low-dimensional structure for populations of neurons that cover a single brain area. As a result, circuit-based models can now capture the genesis of population-wide patterned dynamics. However, we have scant results about how low-dimensional dynamics affect the interaction between distributed, but connected, brain regions. Using a layered spiking network with within- and between-layer spatially structured connectivity, we show that the activity state of a downstream area influences the quality of communication from an asynchronous upstream area. More specifically, if a downstream area is shifted from a stable asynchronous regime to an unstable pattern-forming regime, a previously proposed low-dimensional interaction measure, termed communication subspace, is disrupted. Therefore, we show that intrinsic pattern-formation not only affects within-area dimensionality, but also impairs information transmission from upstream brain areas. These results offer an important mechanistic insight into how neuronal state can affect the flow of information in distributed brain circuits.

2-062. Dynamic allocation of limited memory resources in reinforcement learning

Nisheet Patel[1,2] NISHEET.PATEL@UNIGE.CH
Luigi Acerbi[3] LUIGI.ACERBI@HELSINKI.FI
Alexandre Pouget[1] ALEXANDRE.POUGET@UNIGE.CH

[1]University of Geneva
[2]Department of Basic Neuroscience

[3]University of Helsinki

Biological brains are inherently limited in their capacity to process and store information, and yet are capable of solving complex tasks with apparent ease. Intelligent behavior is related to these limitations since resource constraints drive the need to generalize and assign importance differentially to features in the environment or memories of past experiences. Recently, there have been parallel efforts in reinforcement learning and neuroscience to understand strategies adopted by artificial and biological agents to circumvent limitations in information storage. However, the two threads have been largely separate.

In this work, we propose a dynamical framework to maximize expected reward under constraints of limited resources, which we implement with a cost function that penalizes precise representations of action-values in memory, each of which may vary in its precision. We derive from first principles an algorithm, Dynamic Resource Allocator, which we apply to two standard tasks in reinforcement learning and a model-based planning task, and find that it allocates more resources to items in memory that have a higher impact on cumulative rewards. Our work thus provides a normative solution to the problem of learning how to allocate costly resources to a collection of uncertain memories.

Our work also partly explains a commonly reported phenomenon in neuroscience: while early sensory and somatosensory brain areas recruit more neurons over the course of training, other brain areas responsible for higher-level cognitive processes such as accessing and storing memories (prefrontal cortex) and evidence integration during decision-making (posterior parietal cortex), either commit more neurons or show higher levels of activity during earlier stages of learning than later stages when animals are well-trained. By showing that resource-constrained agents can accelerate learning by starting with more resources, we provide a plausible hypothesis for this observation.

2-063. Modeling neural variability in deep networks with dropout

Xu Pan[1,2] PANXU001@GMAIL.COM
Ruben Coen-Cagli[3] RUBEN.COEN-CAGLI@EINSTEINMED.ORG
Odelia Schwartz[1] ODELIA@CS.MIAMI.EDU

[1]University of Miami
[2]Computer Science
[3]Albert Einstein College Of Medicine

Recent studies show that deep Convolutional Neural Network (CNN) models successfully predict neural responses to visual inputs. However, those studies only focused on average neural activation. Unlike a deterministic system, neural responses have high trial-by-trial variability, which has important consequences on the amount of information that is encoded. We ask if CNN models can also capture key properties of cortical neural variability. To address this question, we consider dropout, a regularization technique that sets random subsets of units to zeros at each training mini-batch. This study builds on a previous theoretical work that showed dropout can be reinterpreted as sampling latent variables (network weights) under a Bayesian statistics framework, and used to estimate uncertainty. We train CNNs on object classification in natural images, and find that response variability of the hidden units induced by test-time dropout, qualitatively reproduces a wide range of experimental observations from visual cortex. Consistent with the data, population variability in our model is low-dimensional, and the amount of dropout affects dimensionality in a similar way as visual attention does in cortex. Our model displays small, positive noise correlations that are positively related to filter similarity (limited-range). Furthermore, noise correlations induced by dropout tend to decrease during training and, after training, they limit decoding accuracy. Interestingly, we find that classification accuracy can be further improved without additional training, only by removing the correlations at test time (trial-shuffling). To our knowledge, this study is the first modeling work that shows how noise correlations could emerge in deep CNNs from a sampling-based representation of uncertainty, and quantifies their impact on a visual task.

2-064. A functional role of cortical oscillations for probabilistic computation in spiking neural networks

Michael G. Muller[1,2] MUELLER@IGI.TUGRAZ.AT
Agnes Korcsak-Gorzo[3] A.KORCSAK-GORZO@FZ-JUELICH.DE
Walter Senn[4] WALTER.SENN@UNIBE.CH
Luziwei Leng[5] LUZIWEI.LENG@KIP.UNI-HEIDELBERG.DE
Andreas Baumbach[5] ANDREAS.BAUMBACH@KIP.UNI-HEIDELBERG.DE
Oliver Julien Breitwieser[3] OLIVER.BREITWIESER@KIP.UNI-HEIDELBERG.DE
Sacha J. van Albada[6] S.VAN.ALBADA@FZ-JUELICH.DE
Karlheinz Meier[3] MEIERK@KIP.UNI-HEIDELBERG.DE
Robert Legenstein[1] ROBERT.LEGENSTEIN@IGI.TUGRAZ.AT
Mihai Petrovici[7,8] MIHAI.PETROVICI@UNIBE.CH

[1]Graz University of Technology
[2]Institute of Theoretical Computer Science
[3]Heidelberg University
[4]Universitat Bern
[5]Heidelberg University and University of Bern
[6]University of Cologne and Julich Research Centre
[7]University of Bern and Heidelberg University
[8]Department of Physiology and Kirchhoff-Institute for Physics

There is mounting evidence that the brain deals with uncertainty by carrying out probabilistic inference. This adds support to the long-standing idea that the brain employs a generative model to capture regularities in the world in the form of probability distributions, with neural dynamics representing a form of sampling from these distributions. However, real-world distributions are highly complex and pronouncedly multimodal, rendering computations such as sampling difficult in practice. We propose that this problem is overcome through background oscillations, a ubiquitous phenomenon throughout the brain, which we suggest to effectively implement a form of tempering in networks of spiking neurons. We show how the intensity of background input defines a Boltzmann temperature for spiking neurons. This allows oscillating background input to induce a tempering schedule and promote different modes of operation in spiking neural networks, resulting in improved coverage of the relevant state space. High levels of background input lead to periods of high temperature, where the state space is rapidly traversed, while periods with low input result in low temperatures, where the network state converges to a single mode. While the resulting impact on the network's probability distribution can be analytically described for current-based neurons, we show that this tempering effect is also present in conductance-based models across a broad range of physiologically relevant parameter settings. In particular, we demonstrate how background oscillations improve mixing in a stimulus interpretation task with ambiguous input, a scenario highlighting the advantages of sampling models. These networks show the emergence of a phase-based encoding, with coherent interpretations having a high probability at certain parts of the oscillation cycle. Our work thus provides a rigorous framework for the suggested functional role of cortical oscillations as a tempering mechanism and creates a novel link from sampling-based computations in spiking neural networks to neural coding.

2-065. Cholinergic-mediated adaptive learning in cortical microcircuits

Heng Wei Zhu[1,2] HENGWEI.ZHU@BRISTOL.AC.UK
Will Greedy[1,3] WILL.GREEDY@BRISTOL.AC.UK
Jack Mellor[1] JACK.MELLOR@BRISTOL.AC.UK
Rui Ponte Costa[1] RUI.COSTA@BRISTOL.AC.UK

[1]University of Bristol
[2]School of Physiology, Pharmacology and Neuroscience
[3]Computer Science

Adaptive learning rules are commonly used in deep learning, offering improvements to learning speed, generalisation and robustness to noise. Such properties are a hallmark of biological neural networks, but it is not known whether the brain could be using similar adaptive learning rules. Here, we introduce a model of the cholinergic system as an adaptive learning system that modulates synaptic updates derived from a backprop-like credit assignment mechanism in cortical microcircuits (ChoCA). In this model, the cholinergic system integrates local prediction error signals to modulate credit assignment across a dendritic cortical network of pyramidal cells and somatostatin interneurons. We incorporated different adaptive methods (Adagrad and RMSprop) from machine learning into ChoCA and importantly, tested different levels of cholinergic modulation (neuron, layer and network), consistent with experimental observations, on standard pattern recognition tasks (MNIST and F-MNIST). Our results show that modulation at the level of neurons or layers achieves similar learning performance compared

to standard adaptive methods which require synapse-specific modulation. This demonstrates that the diffuse nature of cholinergic neuromodulation may be sufficient for effective adaptive learning. Furthermore, we show that cholinergic modulation of somatostatin interneurons, in addition to pyramidal cells, is important for learning. Overall, our work demonstrates the efficacy of coupling neuromodulatory signals with efficient credit assignment for adaptive learning in the brain.

2-066. Temporal computations through dynamics on neural manifolds

Manuel Beiran[1,2] MANUEL.BEIRAN@ENS.FR
Mehrdad Jazayeri[3] MJAZ@MIT.EDU
Srdjan Ostojic[1] SRDJAN.OSTOJIC@ENS.FR

[1]Ecole Normale Superieure
[2]Department of Cognitive Studies
[3]Massachusetts Institute of Technology

Animals process temporal information flexibly to control the timing and speed of their actions depending on behavioral requirements. Recordings in monkeys performing timing tasks have shown that neural activity at the population level evolves along low-dimensional manifolds where the speed of neural trajectories can be flexibly controlled. Here we examine how such manifolds are generated within recurrent network models, and how their properties subserve temporal computations.

Based on experimental observations and trained RNNs, we identified four necessary features of neural manifolds required to perform such computations: (F1) the recurrent connectivity generates attractive manifolds in neural space, (F2) the dynamics along the manifold are shaped by the connectivity, (F3) the speed along the manifold can be modulated by external inputs, and (F4) the same network can generate multiple manifolds to account for different task epochs. Secondly, we develop minimal, analytically tractable models of recurrent networks with low-rank connectivity to generate such manifolds, linking specific connectivity properties to each of the required dynamic features. Then, we exploited the network models to implement a set of timing tasks.

We show that recurrent networks with similar correlations within the low-rank connectivity structure are sufficient to generate slow manifolds (F1). Deviations from a perfect isotropic connectivity robustly shape the dynamics along the slow manifold (F2), while tonic inputs correlated with the connectivity patterns can modulate the speed (F3). Different manifolds can be generated simultaneously by adding different low-rank connectivity structures (F4). Finally, by mapping the input vectors in timing tasks to the connectivity, these minimal network models serve as the basis to implement a wide range of temporal tasks such as inferring temporal intervals between stimuli, storing intervals in memory or executing an action at different time intervals. Altogether, the models we introduced provide a mechanistic framework for understanding the role of neural manifolds in flexible temporal computations.

2-067. Dynamic predictive coding of natural movies produces visual cortical space-time receptive fields

Linxing Preston Jiang[1,2] PRESTONJ@CS.WASHINGTON.EDU
Dimitrios C. Gklezakos[1] GKLEZD@CS.WASHINGTON.EDU
Rajesh P. N. Rao[1] RAO@CS.WASHINGTON.EDU

[1]University of Washington
[2]Paul G. Allen School of Computer Science & Engineering

The traditional predictive coding model of Rao & Ballard [1] focused on spatial prediction to explain spatial receptive fields and contextual effects in the visual cortex. Here, we introduce a new dynamic predictive coding model that achieves spatiotemporal prediction of complex natural image sequences using time-varying transition matrices. We overcome the limitations of static linear transition models (as in, e.g., Kalman filters) using a recurrent modulation network to adjust the transition matrix dynamically for every time step, allowing the model to predict using a time-varying mixture of possible transition dynamics. At each time step, the network makes a prediction of the next input and estimates a sparse neural code for the input by minimizing prediction error. Minimization of prediction error also allows the network to learn compositional spatial and temporal features. When exposed to natural movies, the model learned localized, oriented spatial filters as well as both separable and inseparable (direction-selective) space-time receptive fields similar to those found in primary visual cortex (V1). Our results suggest that the space-time receptive fields of V1 neurons could be the result of the cortex learning an internal model of the visual world based on predicting its inputs using a sparse dynamic neural code.

2-068. A perturbative approach uncovers context-dependent processing of natural scenes in the retina

Baptiste Lefebvre[1] BAPTISTE.LEFEBVRE.BL@GMAIL.COM
Alexander S. Ecker[2] ECKER@CS.UNI-GOETTINGEN.DE
Thierry Mora[1] THIERRY.MORA@ENS.FR
Ulisse Ferrari[3] ULISSE.FERRARI@INSERM.FR
Olivier Marre[4] OLIVIER.MARRE@INSERM.FR

[1]Ecole Normale Superieure
[2]University of Goettingen
[3]Institut de la Vision
[4]Institut de la Vision, Sorbonne Universite

Neurons in the early visual system are supposed to extract specific features from natural scenes. However, many works have shown that this processing depends on the visual context. It is thus not clear if they always extract the same feature, and therefore a stable representation of the visual world, across different contexts. In the retinal output, some types of ganglion cells respond selectively to light increase (On cells), while others respond to light decrease (Off cells). However, it is unclear if this selectivity, characterized with synthetic stimuli, is maintained when they are stimulated with natural scenes. Here we address this issue by recording the responses of ganglion cells to stimuli composed of natural images slightly perturbed by patterns of random noise. We measured a receptive field from the responses to many checkerboard patterns superimposed on the same natural image, i.e. an image-dependent receptive field. For the same cell, these receptive fields strongly depended on the natural image. When changing the image, the receptive field could switch its polarity from On to Off. Standard, linear models could not reproduce this result. However, a convolutional network model learned on the responses to many natural images could predict these context dependent receptive fields, and suggest that it is the result of the convergence of several circuits onto the same ganglion cell. This strong dependence of the On-Off polarity on the context makes it challenging to decode luminance from these ganglion cells. However, more abstract features could be reliably decoded. This perturbative approach thus uncovered the selectivity of retinal ganglion cells to more abstract features, and could be used in other sensory areas to characterize the selectivity of neurons during natural scene processing.

2-069. Random background connectivity controls interference in sequential learning

Friedrich Schuessler[1] FRSCHU@POSTEO.DE
Francesca Mastrogiuseppe[2] FRAN.MASTROGIUSEPPE@GMAIL.COM
Srdjan Ostojic[3] SRDJAN.OSTOJIC@ENS.PSL.EU
Omri Barak[4] OMRI.BARAK@GMAIL.COM

[1]Technion - Israel Institute of Technology
[2]Gatsby Computational Neuroscience Unit, UCL
[3]ENS Paris
[4]Technion, Haifa

Consider learning a new tune on an instrument. There will be sequences similar to those in previously learned pieces. Learning may either adapt what was previously learned, or form a new representation. In simplified settings, experimental studies showed that both phenomena take place depending on the task setting and on the individual learner. Statistical similarity between stimuli has been found to be a major controlling factor. Mechanisms underlying individual differences, however, have mostly remained elusive. Here we explore the learning dynamics of recurrent neural networks trained sequentially on multiple tasks. We observe a transition between catastrophic forgetting and retaining representations between tasks depending on network parameters. For a simple example of a fixed point task, we determine how training on a second task either leads to rotation of an existing fixed point or the formation of a new one. We find that the strength of random connectivity mitigates the transition between these two regimes: Networks with larger random weights at the beginning of training retain more information between tasks. We also observe a similar pattern for more complex tasks. In summary, our model gives insight into the mechanisms underlying individual differences in learning.

2-070. Distributional reinforcement learning explains ensemble dopamine responses in habenula lesioned mice

Sandra Alejandra Pinto SROMEROPINTO@G.HARVARD.EDU
Naoshige Uchida UCHIDA@MCB.HARVARD.EDU

Harvard University

In traditional reinforcement learning (RL), an agent predicts the average future rewards to guide actions. A novel algorithm called distributional RL (DRL), predicts the distribution of rewards and improves performance in artificial agents. A recent study (Dabney et al., 2020) provided initial evidence that the dopamine system may be working under DRL, accounting for the reported diversity in reward prediction errors (RPE). A previous study (Tian and Uchida, 2015) showed that habenula lesions impaired specific aspects of average dopamine responses while largely preserving RPEs. Our re-examination of this dataset, however, revealed a change in ensemble responses not explained by traditional RL. Here, we tested whether these changes conform to DRL predictions.

Mice were trained in a task in which odors predicted reward with different probabilities (10, 50 and 90%). In this task, traditional RL predicts that cue responses are proportional to reward probability. This was the case in controls but not in lesioned animals. We hypothesized that DRL provides additional constraints on the ensemble response pattern of cue-evoked responses. We categorized them into levels of optimism based on whether their amplitude in 50% reward trials was above or below the interpolated response from 10% and 90% reward trials. While we found similar proportion of optimistic and pessimistic responses in controls, the distribution was skewed toward optimistic responses after lesions. We tested whether this can be explained by dopamine neurons' asymmetries in reward responses to positive and negative prediction errors. As opposed to traditional RL, a DRL model incorporating reward response asymmetries reproduced the optimistic bias in cue responses. Accordingly, the decoded distribution associated to the 50%-reward cue based on an expectile code was biased to optimistic values after lesions. Together, these results demonstrate that a change in ensemble dopamine responses, not explained by traditional RL, could be readily explained by DRL.

2-071. Parallel inference in the mouse frontal cortex

Fanny Cazettes[1] FANNY.CAZETTES@NEURO.FCHAMPALIMAUD.ORG
Masayoshi Murakami[2] MASAYOSHI.MURAKAMI@NEURO.FCHAMPALIMAUD.ORG
Alfonso Renart[1] ALFONSO.RENART@NEURO.FCHAMPALIMAUD.ORG
Zachary Mainen[1] ZMAINEN@NEURO.FCHAMPALIMAUD.ORG

[1]Champalimaud Centre for the Unknown
[2]Department of Neurophysiology, University of Yamanashi, Japan

Environments supply organisms with streams of raw sensory information that must be processed in order to infer the latent states of the environment relevant to a given situation. It is believed that, while early sensory cortices proceed in parallel, higher order cortices perform complex operations such as hidden state inference by integrating only selected, behaviorally-relevant information. However, here we show that frontal cortical circuits simultaneously represent a complete range of inferred latent variables (LVs) regardless of behavioral relevance. We used Neuropixels probes to record large ensembles of neurons in the frontal cortex while mice performed a probabilistic foraging task. In this task, mice had to use the sequence of successful and failed foraging attempts to infer a particular LV in order to optimally time the duration of their foraging bout. After training, this LV predicted the mice's behavior and could be accurately decoded from frontal cortical neural ensembles. Remarkably, however, linear decoding could also recover a complete basis set of alternative LVs with no relationship to the actual behavior. We showed that the additional LVs were genuinely represented in the frontal cortex and not only by virtue of their correlations with the relevant LV or with movement signals. These results demonstrate that multiple inferences can be performed in parallel and suggest that attentional selection may occur later than generally thought. Parallel inferential processing could permit behavioral adaptation without implementation of new computations, providing a rapid and powerful form of cognitive flexibility.

2-072. Quantifying brain state transition cost via Schrodinger's bridge

Genji Kawakita[1] GKAWAKI1@SWARTHMORE.EDU
Shuntaro Sasai[2] SHUNUNYS@GMAIL.COM
Jun Kitazono[3] JUN.KITAZONO@GMAIL.COM
Masafumi Oizumi[3] C-OIZUMI@G.ECC.U-TOKYO.AC.JP

[1]Swarthmore College
[2]University of Wisconsin-Madison

[3]The University of Tokyo

To effectively solve tasks with varying degrees of cognitive demands, the brain must swiftly transition between brain states. However, measuring the transition cost from the recording of brain activity has remained difficult. We present a framework to quantify the transition cost by reframing brain state transition as a Schrodinger's bridge problem, the problem of finding the most likely transition path between endpoint probability distributions, proposed by Erwin Schrodinger in 1931. Under this framework, we can consider the brain activity as a stochastic process following a baseline transition probability, which may be modulated by some external input. We can obtain the brain state transition cost as the Kullback-Leibler divergence between baseline and controlled (modulated) transition probabilities. To examine the utility of our method, we used fMRI data of resting state and task states from the Human Connectome Project (HCP). We constructed the transition probability matrices and probability distributions of all the states. Setting the resting state transition probability to the baseline, we computed the transition costs to the task states. We found that the transition cost to the more cognitively demanding (difficult) task state was larger than that of the less demanding (easy) task state in the working memory task. In addition, we found that the transition cost from easy to difficult task state was larger than the cost for the reverse transition. However, when we changed the baseline to the transition probabilities of the task states, we did not observe both of the findings. These results suggest that our method might capture the cognitive load of subjects from the recording of their brain activities, and are consistent with the previous findings of resting state dynamics as the

2-073. A theory for Hebbian plasticity in recurrent E-I networks

Samuel Eckmann[1] EC.SAM@OUTLOOK.COM
Julijana Gjorgjieva[1,2] GJORGJIEVA@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]Computation in Neural Circuits Group

The Stabilized Supralinear Network is a model of recurrently connected excitatory (E) and inhibitory (I) neurons with a supralinear input-output relation. It can explain cortical computations such as response normalization and inhibitory stabilization. However, the network's connectivity is designed by hand, based on experimental measurements. How the recurrent synaptic weights can be learned from the sensory input in a biologically plausible way is unknown. Earlier theoretical work on plasticity focused on single neurons and the balance of excitation and inhibition but did not consider the simultaneous plasticity of recurrent synapses and the formation of receptive fields. Here we present a recurrent E-I network model where all synaptic connections are simultaneously plastic, and excitatory neurons self-stabilize by recruiting co-tuned inhibition. Motivated by experimental results, we employ a local Hebbian plasticity rule with multiplicative normalization for excitatory and inhibitory synapses. We develop a theoretical framework that explains how plasticity enables inhibition balanced excitatory receptive fields that match experimental results. We show analytically that sufficiently strong inhibition allows neurons' receptive fields to decorrelate and distribute themselves across the stimulus space. For strong recurrent excitation, the network becomes stabilized by inhibition, which prevents unconstrained self-excitation. In this regime, external inputs integrate sublinearly. As in the Stabilized Supralinear Network, this results in response normalization and winner-takes-all dynamics: when two competing stimuli are presented, the network response is dominated by the stronger stimulus while the weaker stimulus is suppressed. In summary, we present a biologically plausible theoretical framework to model plasticity in fully plastic recurrent E-I networks. While the connectivity is derived from the sensory input statistics, the circuit performs meaningful computations. Our work provides a mathematical framework of plasticity in recurrent networks, which has previously only been studied numerically and can serve as the basis for a new generation of brain-inspired unsupervised machine learning algorithms.

2-074. Upstream spikes phase-locking to gamma oscillations enhances downstream neurons population coding

Tal Dalal[1,2] TALDALAL123@GMAIL.COM
Rafi Haddad[1] RAFIHADDAD@GMAIL.COM

[1]Bar-Ilan University
[2]Neuroscience

Gamma-oscillations are observed in many brain regions. Previous studies have related the increase in gamma-oscillations to improvement in a range of cognitive functions including memory, attention, motivation, cognitive flexibility and sensory responses. Disruption of oscillations was directly related to impairment in olfactory behavior and indirectly to several brain disorders such as Alzheimer's disease, Parkinson's disease, schizophrenia, autism, and epilepsy. However, the neural mechanism subserving these functions is still hotly debated. One

leading hypothesis is that gamma oscillations are involved in synchronizing output neuron spiking, which thus improves information transfer in the brain. Odor stimuli are well known to evoke gamma-oscillations in the olfactory bulb, which require the activity of the GABAergic Granule cells (GCs) that form inhibitory connections with olfactory bulb excitatory output neurons – the Mitral/Tufted cells (MTCs). Here we show that optogenetically increasing the activity of GCs surprisingly enhances downstream neurons odor responses and ensemble representation, despite reducing the spiking activity of the MTCs. These enhancements were not observed when we directly reduced MTCs spiking activity. Furthermore, optogenetically suppressing the activity of GCs, which slightly increased MTCs activity, reduced downstream neuron odor responses and degraded their ensemble representation. Interestingly, optogenetically increasing GCs-mediated inhibition increased MTCs spikes synchronization to the odor-evoked gamma oscillations. This tighter synchronization of spikes to the gamma cycles did not depend on the gamma oscillation power. Our results provide the first evidence for the hypothesis that gamma oscillation enhances information processing in downstream targets, by serving as an internal clock which upstream neuron spikes are synchronized to.

2-075. Multiple hypothesis-dependent sensory and spatial computations in mouse Retrosplenial Cortex

Jakob Voigts[1,2] JVOIGTS@GMAIL.COM
Ingmar Kanitscheider[3] INGMAR@OPENAI.COM
Nicholas J. Miller[1] NICJMILL@MIT.EDU
Enrique H.S. Toloza[1] QTOLOZA@MIT.EDU
Jonathan P. Newman[1] JPNEWMAN@MIT.EDU
Ila R. Fiete[1] FIETE@MIT.EDU
Mark T. Harnett[1] HARNETT@MIT.EDU

[1]Massachusetts Institute of Technology
[2]McGovern Institute for Brain Research at MIT
[3]OpenAI

Cognitive function requires holding multiple hypotheses in mind and interpreting sensory information differently depending on these hypotheses. Here, we examined the cortical substrates of this process in a novel spatial inference task through modeling and electrophysiology in mouse retrosplenial cortex (RSC). The task involves self-localization from an unknown starting position in a familiar dark arena containing two identical visual landmarks in a fixed spatial relationship. After encountering one of the landmarks, an ideal observer's uncertainty would shrink to two potential location hypotheses. After encountering the second landmark, the observer would be able to resolve the remaining uncertainty and accurately self-localize. We trained a recurrent artificial neural network (ANN) on the task to predict the latent variables and population-level neural representations necessary to solve the task. The predictions included a representation of uncertainty in which the same location hypothesis is encoded differently depending on the level of uncertainty, as well as hypothesis-dependent sensory processing. We then trained mice to solve this complex memory-dependent multi-hypothesis probabilistic localization task. Chronic tetrode recordings of neurons in RSC revealed population dynamics corresponding to multi-hypothesis representation and computation. We observed that RSC neurons encode a mouse's possible location, while also retaining information about the remaining uncertainty. This encoding was conjunctive: individual neurons code for a location as well as a level of uncertainty, self-motion, and other variables. Additionally, the same visual landmark information is processed differently in RSC depending on the current hypothesis. Our findings establish that RSC encodes probabilistic spatial information, shaping how data is integrated based on prior knowledge of the environment and ongoing evolving hypotheses. Thus, mouse associative cortex holds mixed representations of hypotheses and sensory data and performs flexible context-dependent computations. Our findings provide a functional framework for understanding how associative cortex instantiates cognitive operations like hypothesis creation and selection.

2-076. Optimal allocation of finite sampling time in accumulator models of multialternative decision making

Jorge Ramirez-Ruiz[1,2] JORGE.RAMIREZ@UPF.EDU
Ruben Moreno Bote[3,4] RUBEN.MORENO@UPF.EDU

[1]Universitat Pompeu Fabra
[2]Departament de Tecnologies de la Informacio i les Comunicacions
[3]University Pompeu Fabra, Barcelona
[4]Center for Brain and Cognition

In multialternative decision making we can sample many alternatives, but very rarely do so: often we ignore arbitrarily many of the options and focus only on a bunch of them. Although behaviors like this can be a sign of heuristics, they can also correspond to optimal behavior under limited resources. Here we study the problem of how to allocate limited sampling capacity to multiple alternatives, modelled as accumulators of noisy evidence over time. The decision maker first allocates finite sampling time arbitrarily to the accumulators, and then receives feedback and reward equal to the highest inferred rate of accumulation. We find that this model undergoes an optimal policy transition as a function of sampling capacity. For small capacity, the optimal number of accumulators to sample is constant and equal to five, regardless of the prior, thus ignoring all other alternatives. For large capacities, the optimal number of sampled accumulators grows close to a power law for Gaussian, uniform and bimodal priors. We found that allocating equal time to the sampled accumulators is better than using uneven time allocations. Our work highlights that multialternative decisions are endowed with breadth-depth tradeoffs, demonstrates how the optimal balance depends on the amount of limited resources, and shows that a small number of considered alternatives is always optimal for small capacities.

2-077. Context-dependent population dynamics in auditory cortex during flexible behaviour

Joao Barbosa[1,2] PALERMA@GMAIL.COM
Remi Proville[3] REMI@TAILOREDDATASOLUTIONS.FR
Chris Rodgers[4] XRODGERS@GMAIL.COM
Srdjan Ostojic[1] SRDJAN.OSTOJIC@ENS.FR
Yves Boubenec[1,5] YVES.BOUBENEC@ENS.FR

[1]Ecole Normale Superieure
[2]Group for Neural Theory
[3]Tailored Data Solutions, SAS, Cavaillon, France
[4]Department of neuroscience, Columbia University
[5]Cognitive Studies

Understanding how stable neural circuits implement flexible, context-dependent behavior is crucial to understand cognition. The classical view posits that sensory areas do not play a computational role beyond the representation of sensory stimuli, while downstream areas apply context-dependent rules to guide behavior [de Lafuente & Romo (2006)]. On the other hand, recent empirical evidence has shown that primary sensory areas are modulated by behavioural context [Otazu et al. (2009); Niell & Stryker(2010)] and that the information encoded in these areas is task dependent [Fritz et al. (2003)]. How the brain selects relevant stimuli and ignores irrelevant ones – and in particular the computational role of sensory areas in this selection – remains to be fully elucidated. Here, we focus on the population dynamics previously recorded from the primary auditory cortex(A1), while rats performed a task that required flexible, context-dependent behavior [Rodgers & deWeese(2014)]. Specifically, rats performed stimuli discrimination (go/no-go) according to either the location or the pitch of the auditory stimuli, depending on the context that was not explicitly cued. Based on recurrent neural networks (RNNs) trained on context-dependent tasks, we developed a minimal RNN model that implements the specific go/no-go task of [Rodgers & deWeese (2014)]. While previous population analyses on the same dataset reported indistinguishable representations for selected and ignored stimuli, our model predicts that different populations selectively gate the integration of different stimuli, in a context-dependent fashion. Motivated by these predictions, we identified 2 distinct populations in A1, based on their context-specific activity during the pre-stimulus period. We found that, on the collective level, each population discriminates more strongly the relevant stimuli inits preferred context. Our results reveal clear context-dependent dynamics in A1 and suggest it plays a central role in flexible computations.

2-078. Division versus subtraction: task-specific gating in recurrent neural networks

Milton Montero[1,2] M.LLERAMONTERO@BRISTOL.AC.UK
Abdulah Khwaja[1] MK17486@BRISTOL.AC.UK
Rui Ponte Costa[1] RUI.COSTA@BRISTOL.AC.UK

[1]University of Bristol
[2]Computational Neuroscience Unit

Inhibition gates information flow in neural circuits. There is evidence for both divisive and subtractive inhibition-mediated gating. However, why does the brain sometimes prefers one form of gating over the other has remained unclear. Here, we highlight the benefits of these two modes of gating for learning using simple gated-RNNs. First, we show analytically that divisive gating amplifies input variability more than subtractive gating. This suggests that subtractive-RNNs should yield more robust learning dynamics, which we tested using a combination of working memory and pattern recognition tasks. Our experiments show that learning with sub-RNNs is generally more robust (i.e. less noisy), consistent with the theoretical predictions. Next, using a dimensionality reduction analysis we show that sub-RNNs tend to yield less diverse solutions, consistent with the less noisy learning dynamics. Finally, we demonstrate that divisive gating performs well when in tasks with discrete gating while sub-RNNS perform better in tasks that encourage continuous gating. Overall, our results suggest that the specific form of gating employed by the brain should depend on the task and input statistics faced by a particular neural circuit.

2-079. Odor-evoked increases in spiking variability in the olfactory bulb

Andrea Barreiro[1,2] ABARREIRO@SMU.EDU
Cheng Ly[3] CLY@VCU.EDU
Shree Hari Gautam[4] GAUTAM.SHREEHARI@GMAIL.COM
Woodrow Shew[4] WOODROWSHEW@GMAIL.COM

[1]Southern Methodist University
[2]Department of Mathematics
[3]Virginia Commonwealth University
[4]University of Arkansas

Neurons in the olfactory bulb (OB) process odor information and relay it to other brain regions via coordinated spiking activity. A general principle of cortical sensory processing is that the onset of sensory input results in a decrease in spiking variability and covariability across neurons. Since reduced variability and covariability is thought to be beneficial for sensory coding, one might guess that this principle applies to all sensory circuits, including OB, but this possibility remains untested. Here we recorded spiking activity of many single neurons in OB during olfactory stimulation and found that variability and covariability increased, rather than decreased, at stimulus onset. To gain insight into this finding we analyzed a firing rate model of OB, focusing on the possibly synergistic effects of two features known to individually modulate mitral cell (MC) spike dynamics: inhibitory interactions within OB and variability of input from olfactory receptor neurons (ORNs). We found that increases in the variance of ORN inputs were not needed to explain the increase in OB spiking variance we observed in our experiments. We found that uncorrelated granule cell inhibition must be strong, but shared granule cell inhibition must be weak. Our findings establish increased spiking variance and covariance at the onset of input as a viable alternative coding strategy to the more commonly observed decrease in variability found in many cortical systems.

2-080. Circuit analysis of midbrain inhibitory and excitatory populations during decision maintenance

Yuxin Pan[1,2] YUXINPAN@ION.AC.CN
C. Ann Duan[1,3] ANNDUAN2@GMAIL.COM
Guofen Ma[4] GUOFENMA@SJTU.EDU.CN
Siyu Zhang[4] ZHANG_SIYU@SJTU.EDU.CN
Ninglong Xu[1] XUNL@ION.AC.CN

[1]Chinese Academy of Sciences
[2]Institute of Neuroscience
[3]New York University Shanghai
[4]Shanghai Jiao Tong University School of Medicine

Maintenance of decision-related information bridges past events with future actions, and is essential for survival in a dynamic environment. Although the neural correlates and computational models of decision maintenance have been investigated extensively, specific neural circuit underpinnings of these cognitive processes remain elusive, making it difficult to constrain biologically plausible models. Here, we employ pathway- and cell-type-specific methods to probe neural circuit mechanisms of perceptual decisions in the midbrain superior colliculus (SC), a key subcortical node for cognition and action. Previous work has shown that both SC and its major cortical input, the secondary motor cortex (M2), are essential for memory-guided auditory discrimination in head-fixed mice: M2 sends choice-related information to SC neurons, and this pathway is preferentially required for maintenance of harder decisions. In this study, we examine how such information is further transformed within SC. Two-photon calcium imaging of SC neurons revealed an enrichment of choice-related activity in the SC during decision formation, maintenance and execution, compared to its upstream M2 input. To further dissect the functions of SC sub-networks, we combined transsynaptic viral tracing, immunohistochemistry, and whole-cell recording in SC brain slices to confirm that both SC inhibitory and excitatory neurons receive synaptic inputs from M2,

suggesting their potential involvement in the task. We then conducted fiber photometry recording of genetically identified SC sub-populations during task performance. We found differential temporal dynamics in choice coding between excitatory and inhibitory subpopulations, suggesting a gradual reconfiguration of SC sub-networks from balanced, sub-threshold maintenance of choice information during early decision period to coordinated promotion of contralateral choices before action initiation. Together, our experimental evidence provides constraints for decision-making models where competitive inhibition between choices can be achieved by genetically- and functionally-identified midbrain sub-networks.

2-081. Identifying the central circuits for olfactory navigation in walking Drosophila

Andrew MM Matheson[1,2] AMM1420@NYU.EDU
Aaron J Lanz[1] AJL787@NYU.EDU
Timothy A Currier[1] CURRIER@STANFORD.EDU
Angela M Licata[1] ANGELA.M.LICATA@GMAIL.COM
Katherine I Nagel[1] KATHERINE.NAGEL@NYULANGONE.ORG

[1]NYU Langone School of Medicine
[2]Neuroscience Institute

All organisms must navigate to find food. Many navigate in turbulent chemical environments, where odour alone is insufficient to locate food sources. A common algorithm in such environments involves surging upwind during odour, and casting crosswind at odour offset. Previously, we developed a behavioural paradigm for walking Drosophila, where flies run upwind during appetitive odour, and search locally at odour offset. Here, we used this paradigm to identify circuits that translate sensory stimuli into navigation behaviours.

Using a combination of optogenetic activation, and 2-photon imaging we mapped olfactory circuitry from the periphery, through the lateral horn (LH), and mushroom body (MB), to the fan-shaped body (FSB). At each stage we identified multiple groups of odour-responsive neurons that promote navigation phenotypes including upwind orientation and offset search, when activated. In the LH and MB, we find neurons encode odour non-directionally and drive complete navigation programs (upwind run + offset search). Within the FSB we observe neurons whose responses depend on odour and wind-direction, and activation elicits single navigation components (upwind orientation or turning). We hypothesize the FSB plays a key role in translating sensory input into navigational output.

To understand how the FSB might perform this role, we are combining our activation and physiological data with analysis of the hemibrain connectome. Our data suggests the FSB receives pure wind and pure odour inputs. Wind-sensitive neurons output to local neurons, which couple asymmetrically to left and right projecting output neurons. Intriguingly, local neurons that receive olfactory input have different output asymmetries than those without. We hypothesize olfactory inputs gate the coupling of wind-direction to outputs with different left-right asymmetries as inputs target the output tufts of local neurons. This may allow flies to flexibly change their wind orientation depending on odour input. We are developing a computational model to test these ideas.

2-082. Neural network model of amygdalar memory engram formation and function

Rylan Schaeffer[1] RYLANSCHAEFFER@G.HARVARD.EDU
Nimrod Shaham[1] NSHAHAM@G.HARVARD.EDU
Gabriel Kreiman[1] GABRIEL.KREIMAN@TCH.HARVARD.EDU
Haim Sompolinsky[2] HAIM@FIZ.HUJI.AC.IL

[1]Harvard University
[2]Hebrew University; Harvard University

The past decade has seen an explosion of experimental research into the formation and function of so-called *memory engrams*, distributed sub-populations of neurons that undergo enduring changes in response to learning and that are both sufficient and necessary to retrieve a learnt association. However, theoretical understanding has lagged behind experimental work, leaving important questions unanswered, and a model of memory engrams that explain what computational principles determine how neurons collectively coordinate to encode an experience in an engram or how engrams interact with one another to control behavior is lacking. Here, we provide a novel interpretation connecting engrams in associative learning to a Bayesian behavioral account of learning. We then implement our insights in a rate-based neural network model and use it to explain key experimental findings including (1) how extinguishing one conditioned stimulus can affect responses to an unrelated conditioned stimulus, and (2) when and how unlearning or new learning occurs if reward contingencies change. Our model relies on key circuit-level findings, specifically excitability, competition, and dopamine-controlled synaptic plasticity. Our model is minimal and consistent with both neuroanatomy and biologically plausible learning mechanisms.

2-083. A diverse task-driven characterization of early and mid-level representations of the primate ventral stream

Santiago A. Cadena[1] SA.CADENA721@GMAIL.COM
Konstantin Willeke[1] KONSTANTIN-FRIEDRICH.WILLEKE@UNI-TUEBINGEN.DE
Kelli Restivo[2] KELLI.RESTIVO@BCM.EDU
George Denfield[2] GHDENFIE@BCM.EDU
Edgar Walker[1] EYWALKER@BCM.EDU
Matthias Bethge[1] MATTHIAS.BETHGE@UNI-TUEBINGEN.DE
Fabian H. Sinz[3] FABIAN.SINZ@UNI-TUEBINGEN.DE
Andreas Tolias[2] ASTOLIAS@BCM.EDU
Alexander S. Ecker[4] ECKER@CS.UNI-GOETTINGEN.DE

[1]University of Tubingen
[2]Baylor College of Medicine
[3]University of Tuebingen
[4]University of Goettingen

The goal of normative modeling in sensory systems is to find the computational objectives that sensory brain areas are optimized to solve. Recently, task-driven models have gained popularity as normative models of visual areas of the primate ventral stream as they extract features that offer unprecedented predictive performance. Currently, these approaches primarily use networks trained on object recognition tasks, or more recently, self-supervised tasks that yield equally good features for object categorization. Their success is further supported by their hierarchical correspondence to brain visual areas: V1 is predicted by early layers while higher layers predict best V4 and IT. However, we have few comparisons with diverse tasks that the visual system solves that might work together to facilitate object recognition, or capture other aspects of visual perception, which could reveal novel functional properties and organizations. Here, we evaluate how well the features from networks pretrained on diverse vision tasks explain novel recordings of neurons from V1 and V4 of macaques in response to natural images. We consider 23 networks trained on different objectives from the “taskonomy" project (Zamir, 2018). Based on their representations, these networks were previously organized into “Semantic", “3D", “2D", and “Geometric" task clusters. We found that 1) the differences in performance across tasks is small in V1, but much larger in V4, highlighting the generality of V1 low-level features. 2) V4 is consistently better predicted by semantic, followed by 3D tasks, suggesting that semantic specialization occurs already at this stage. 3) Semantic tasks provide a hierarchical ordering of V1 and V4 areas; but 2D, and geometric tasks fail to do so, underscoring the strong task-dependent characterization of the nonlinear degree of these areas. Overall, our results suggest that diverse task-driven models offer effective ways to study the role of visual cortical areas on relevant behavioral tasks.

2-084. Shaping of cortical adaptation through thalamic synchrony and feed-forward inhibition

Nathaniel Wright[1,2] NWRIGHT40@GATECH.EDU
Peter Borden[1] PETEYBORDEN@GMAIL.COM
Yi Liew[1] YLIEW6@GATECH.EDU
Michael Bolus[1] MFBOLUS@GATECH.EDU
Garrett Stanley[1] GARRETT.STANLEY@BME.GATECH.EDU

[1]Georgia Institute of Technology
[2]Biomedical Engineering

Sensory signals propagate from the periphery to primary sensory cortex, yielding cortical sensory representations that shape perception. The thalamocortical junction is a complex and poorly-understood station in this pathway. Thalamic activity levels are not trivially relayed to cortex, as signal transformation is shaped by properties of feedforward thalamocortical and recurrent intracortical connectivity. Specifically, a given thalamorecipient cortical neuron is sensitive to both the rate and timing of thalamic spikes due to properties of thalamocortical connectivity, but is also primarily driven by other excitatory and inhibitory cortical neurons. This is relevant for sensory processing in noisy sensory environments, as persistent stimulation alters thalamic spike rates and timing, and can depress synapses. This qualitatively predicts cortical adaptation consistent in nature with reports of perceptual adaptation observed across all sensory modalities, but cortical adaptation and its underlying mechanisms are largely unexplored during wakefulness. To address this knowledge gap, we constructed a biophysical network model of excitatory and inhibitory leaky integrate-and-fire neurons subject to inputs motivated by awake thalamic recordings. Our model predicts that a modest decrease in synchronous thalamic spikes results in profound adaptation of cortical excitatory neurons. Due to differential thalamocortical connectivity, inhibitory neurons are more modestly adapted, and the resulting feedforward inhibition helps dampen excitatory firing in the adapted condi-

tion. To test these predictions, we recorded from and optogenetically manipulated VPm thalamus and S1 cortex in the awake mouse whisker pathway during sensory adaptation. These experiments confirmed our predictions, demonstrating profound adaptation of S1 putative excitatory neurons, modest adaptation of putative inhibitory neurons, with a negligible contribution from thalamocortical synaptic depression. Together, these results identify critical roles for spike timing and feedforward inhibition in the transformation of signals across brain areas. Further, they specifically implicate these mechanisms in cortical sensory adaptation, which may in turn shape perception in noisy sensory environments.

2-085. Geometrically aligned attractor manifolds support rapid and reliable remapping between independent network representations of an unchanging environment

Isabel Low[1,2] ILOW@STANFORD.EDU
Alex Williams[1] AHWILLIA@STANFORD.EDU
Malcolm Campbell[3] MALCGCAMP@GMAIL.COM
Scott Linderman[1] SCOTT.LINDERMAN@STANFORD.EDU
Lisa Giocomo[1] GIOCOMO@STANFORD.EDU

[1]Stanford University
[2]Neurobiology
[3]Harvard University

Recent experiments and theoretical models suggest that, under different internal conditions, a single set of sensory features may drive unique neural representations (i.e. “maps”). This type of neural “remapping” could provide the basis for disambiguation of similar sensory experiences into different episodic memories. Existing models of neural remapping have considered how representations of distinct environments could be embedded in a single circuit, leaving open two challenges specific to models of remapping within a single environment. First, in an unchanging environment, neural activity must respond flexibly to changing internal state while reliably preserving sensory information. Second, a model of remapping in a fixed environment must account for map transitions not initiated by changing sensory input. We consider these challenges in the context of 1D circular track environments, represented by 1D ring attractor networks. We predict that geometrically aligned attractor rings, each associated with a different map, would allow positional information to be preserved across remaps. We combine ring attractor and winner-take-all connectivity patterns with small amounts of directed noise to produce rapid transitions between maps. By including “shared neurons” that are active in both maps, we produce the alignment required to preserve positional information. To test our predictions, we obtained large-scale neural recordings in medial entorhinal cortex, which represents navigationally relevant variables and which demonstrates coding flexibility under changing task and behavioral conditions. We found that population activity underwent synchronous transitions between discrete representations of an unchanging environment. We could linearly decode position across remapping events, indicating that positional information was preserved. Further, the attractor manifolds associated with each map were geometrically aligned, as predicted by our model. Altogether, this work advances our understanding of cortical flexibility, providing theoretical and experimental support for a mechanism by which networks could switch between distinct representations of a fixed environment.

2-086. A topological study of point-attractor dynamics of neuronal assemblies in the zebrafish optic tectum

Enrique Hansen[1,2] EHANSEN@BIO.ENS.PSL.EU
Carina Curto[3] CPC16@PSU.EDU
German Sumbre[1] SUMBRE@BIO.ENS.PSL.EU

[1]Institut de Biologie de l’ecole Normale Superieure
[2]Neurosciences
[3]The Pennsylvania State University

The spontaneous dynamics of the optic tectum of the zebrafish larva show a spatiotemporal structure organized in distinct functional neuronal assemblies (groups of neurons containing significantly correlated spontaneous activations). Recent studies [1-3] have shown that the dynamics of these assemblies display characteristics of attractor-like dynamics. These properties are: (i) repetitive visual stimulation activates assemblies that continue being active after the stimulation which indicates reverberation; (ii) they are capable to transition from sparse to full assembly activation; and (iii) they display winner-take-all dynamics. These circuit properties represent an advantageous mechanism for robust visual detection and improve visual acuity.

To further study the neuronal mechanisms underlying the attractor-like properties, we use two-photon light sheet microscopy to capture the dynamics of the entire tectal circuit (~4,000 neurons) with high-temporal resolution in 3D in combination with optogenetic tools, genetic cell type markers (excitatory and inhibitory), and mathematical methods from dynamical systems and topology.

We have detected the assemblies by combining principal component analysis (PCA) and factor analysis for categorization of correlated neurons. By analyzing the neurons' pair-wise correlations, we have found that some neurons belong to only one assembly (unique), whereas others participate in two or more assemblies. We are now trying to correlate this property with cell type (excitatory or inhibitory) to learn about the functional role of the different components of the assemblies. To study the correlation structure of the assemblies, we used Betti curves using the method in [4]. These were obtained from the correlation matrices, and exhibit a stereotypical Betti curve behavior with $\beta 0$ decreasing gradually, and all higher-order Betti number vanishing for all edge densities. The topological analysis of the structure of the assemblies suggests there is an underlying low-rank structure. We propose that this low-rank structure is a signature of the attractor-like dynamics observed in this system.

2-087. Investigating the connectivity and complexity in modes of active motor neuron pools

Martyna Stachaczyk[1,2] MKS316@IC.AC.UK
Irene Mendez Guerra[1] IRENE.MENDEZ17@IMPERIAL.AC.UK
Dario Farina[1] D.FARINA@IMPERIAL.AC.UK

[1]Imperial College London
[2]Bioengineering

Peripheral movement control is established by the activity of motor neurons innervating muscles1. The trains of motor neurons' action potentials (MUAPs) have been previously exploited for pattern recognition2,3; however, there have been no attempts to explore the inter-motor neuron dynamics that could shed the light on the lower-level movement control and movement generation. To fill this gap, we decomposed electromyographic signals into MUAPs, and analysed the neural activity that gives rise to hand movements. Our dataset comprised simple hand movements (i.e. single degree of freedom movements) and complex hand movements (i.e. two or three simple movement components actuated at the same time). We found that the control of simple and complex hand movements is the result of co-activation of motifs of motor neuron pools, which form structural clusters across forearm muscles. The clusters activated during simple movements are particularly consistent over time, as well as movement class-specific. Next, we established that the clusters' patterns arising for complex hand movements are not a result of a linear summation of their simple components, but rather a consequence of the co-activity of the existing 'simple' and novel motor neuron clusters. We also observed that the statistical distances between the neural activity distributions are inversely proportional to the strength of activation of the degrees of freedom, suggesting that the structure of clusters is dynamic and changes with the motor task complexity. Finally, we used information theory to investigate the extent to which complex movements can be explained by their components. In accordance with our analysis of structural clusters, we found that different simple movements reduce uncertainty in complex movements in a non-uniform manner. Altogether, our results show that (1) motor neurons form clusters that are structurally and functionally specialised, and (2) complex movement derives from the interplay between the simple movements clusters and new clusters.

2-088. A local temporal difference code for distributional reinforcement learning

Pablo Tano[1,2] PABLOTANORETAMALES@GMAIL.COM
Peter Dayan[3] PETER.DAYAN@TUEBINGEN.MPG.DE
Alexandre Pouget[1] ALEXANDRE.POUGET@UNIGE.CH

[1]University of Geneva
[2]Department of Basic Neurosciences
[3]MPI for Biological Cybernetics

Recent theoretical and experimental results suggest that the dopamine system implements distributional temporal difference backups, allowing learning of the entire distributions of the long-run values of states rather than just their expected values. However, the distributional codes explored so far rely on a complex imputation step which crucially relies on spatial non-locality: in order to compute reward prediction errors, units must know not only their own state but also the states of the other units. It is far from clear how these steps could be implemented in realistic neural circuits. Here, we introduce the Laplace code: a local temporal difference code for distributional reinforcement learning that is representationally powerful and computationally straightforward. The code decomposes value distributions and prediction errors across three separated dimensions: reward magnitude (related to distributional quantiles), temporal discounting (related to the Laplace transform of future rewards) and time horizon (related to eligibility traces). Besides lending itself to a local learning rule, the decomposition recovers the temporal evolution of the immediate reward distribution, indicating all possible rewards at all future times. This increases representational capacity and allows for temporally-flexible computations that immediately adjust to changing horizons or discount factors.

2-089. Extended differential covariance recovers ground truth connectivity in multi-scale neural systems

Yusi Chen CYUSI@UCSD.EDU
Terrence J Sejnowski TERRY@SALK.EDU

Salk Institute for Biological Studies

Our ability to sense, think, and react emerges from neural interactions at all scales, thus methods investigating such causal interactions should be central to the study of brain functions. A rich repertoire of statistical methods has been introduced to the field. Still, the pairwise covariance matrix and the precision matrix remain popular because of their intuitive calculation and high sensitivity to network connectivity detection [1]. By calculating the covariance between the derivative signal and the original signal, differential covariance [2] detects network interactions in a dynamical system and retains the intuitive nature of covariance methods. It also echos the assumptions of regression DCM [3], while the estimation speed is much faster. Here we refined differential covariance estimators in linear systems and extended it into nonlinear systems. The refined linear estimator ($\triangle$P) and the nonlinear estimator ($\triangle$ReLu) effectively reduced false positive connections in toy models of linear and nonlinear dynamics. As expected, $\triangle$ReLu is especially powerful in nonlinear dynamics. Besides, both $\triangle$P and $\triangle$ReLu outperformed the precision matrix in detecting ground truth connections in a sparsely connected leaky integrate and fire (LIF) network [4] and in an anatomical connection supported brain surface model [5].

2-090. A cortical microcircuit model explains contextual modulation in primary visual cortex

Mario Dipoppa[1] MARIO.DIPOPPA@GMAIL.COM
Andreas J. Keller[2] ANDREASJAKOB.KELLER@UCSF.EDU
Morgane M. Roth[2] MORGANE.ROTH@UCSF.EDU
Alessandro Ingrosso[1] INGROSSORAMA@GMAIL.COM
Matthew S. Caudill[3] MSCAUDILL@GMAIL.COM
Massimo Scanziani[2] MASSIMO@UCSF.EDU
Ken Miller[1] KDM2103@COLUMBIA.EDU

[1]Columbia University
[2]University of California, San Francisco
[3]University of California, San Diego

The saliency of a stimulus is modulated by the context in which it is presented. Contextual modulation manifests in the dependence of visual responses on the relative orientation between a stimulus and its surround. We found in mouse V1 that contextual modulation differs across cell types. In the iso (cross) condition, meaning the stimulus and surround were similar (orthogonal), excitatory, PV and VIP neurons were inactive (active) while SOM neurons were active (responded weakly). The complementary behavior of SOM and VIP neurons and their connectivity suggested that the VIP->SOM->excitatory disinhibitory circuit may regulate contextual modulation in excitatory neurons. To test this hypothesis, we developed a recurrent model incorporating different cell types. Using non-negative least squares regression and then backpropagation through time, we optimized the synaptic strengths to match the neural responses of the model to those measured experimentally while respecting biological constraints (Figure 1). To test whether the disinhibitory circuit was necessary for the difference between iso and cross responses (Figure 2A), we froze the VIP->SOM inputs to their iso level while presenting a cross stimulus (Figure 2B). The final state was very close to the iso responses. To test whether the disinhibitory circuit was sufficient, we froze the VIP->SOM inputs to their cross-level while presenting an iso input (Figure 2C). This brought the final state close to the cross response. Thus, activation of the disinhibitory pathway is both necessary and sufficient to produce the transition from iso to cross responses. The pathway acts by modulating recurrent excitation in a strongly recurrent circuit, not as a simple feedforward pathway. The model predicted that suppressing VIP neurons will reduce contextual modulation in excitatory neurons (Figure 3A). The optogenetic experiment confirmed our predictions (Figure 3B), supporting the hypothesis that contextual modulation is regulated by the VIP/SOM disinhibitory circuit.

2-091. The emergence of a collective threshold in the response of ant colonies to sensory perturbations

Asaf Gal ASAFG1@GMAIL.COM
Daniel Kronauer DKRONAUER@ROCKEFELLER.EDU

The Rockefeller University

The sensory threshold is one of the most fundamental and well-studied computational primitives organisms perform, both as a standalone computation, and as a component of more complex tasks. In social organisms such as bee swarms and ant colonies, which perform computational tasks at the group level, the collective sensory threshold is an emergent property that depends on the responses of individuals in the group and on the interactions between them. Here we study this emergence in the clonal raider ant (Ooceraea biroi), a model system that provides convenient and precise control over the properties of the colony. We use automated individual tracking to show that an ant colony indeed responds collectively to step changes in temperature, and that this response is characterized by a threshold. We further show that this threshold is sensitive to the size of the colony, implying that interactions play an important role in the response dynamics of the colony, and that the collective threshold is indeed an emergent property that differ from the sensory threshold of the individual ants. We then use a mathematical model to study how collective threshold can emerge in an interacting group of agents and show that an asymmetric, change resisting interaction is required to replicate the experimental observation. Finally, we discuss how heterogeneity and variability between individuals in the group affect this emergence both in the model and in the experiment. Inspired by the history of computational neuroscience, we believe that studying simple responses to well controlled stimuli in can advance our understanding of how sophisticated cognitive-like function emerge in complex social groups.

2-092. Neural signatures of supervised learning vs. reinforcement learning in brain-machine interface tasks

Jacob Portes[1,2] JPP2139@COLUMBIA.EDU
James Murray[3] JMURRAY9@UOREGON.EDU

[1]Columbia University
[2]Neurobiology and Behavior
[3]University of Oregon

Motor skill learning is known to involve synaptic plasticity in motor cortex, but the rules governing synaptic plasticity are not well understood. One plausible candidate is supervised learning (SL), in which a multidimensional error signal is minimized using an internal model (i.e. an estimate of how the neural activity relates to movement). Another candidate is model-free reinforcement learning (RL), in which a scalar reward is maximized by reinforcing exploratory noise in neural activity that leads to positive outcomes. Using recently proposed algorithms for implementing biologically plausible versions of SL and RL in recurrent neural networks (RNNs), we show via mathematical arguments and simulations that it is possible to infer the learning rule from observed RNN activity. In addition to the observed neural activity, this inference requires knowledge of the mapping from neural activity onto behavior, and, because this mapping can be defined by the experimenter in brain-machine interface experiments, we propose that BMI experiments afford an ideal opportunity to test whether SL or RL provides a better description of synaptic plasticity in motor cortex.

2-093. Volatility influences exploration in reward-guided decision-making

Cathy Chen[1,2] CHEN5388@UMN.EDU
R. Becket Ebitz[3] REBITZ@GMAIL.COM
Evan Knep[1] KNEPX001@UMN.EDU
Collin S. Meyer[1] MEYE2772@UMN.EDU
Alexander B. Herman[1] HERMA686@UMN.EDU
Nicola M. Grissom[1] NGRISSOM@UMN.EDU

[1]University of Minnesota
[2]Department of Psychology
[3]Universite de Montreal

In an uncertain world, we constantly balance two goals. We exploit rewarding options when they are available, but also explore alternatives that provides better reward or new information about the world. One factor that seems to

influence exploration is the volatility of the environment, with more exploration in more volatile environments. This makes sense because there is little information to be gained from exploring alternatives in a stable environment, whereas uncertainty about alternative options grows quickly in a volatile environment, as does the potential information gain from exploration. However, in overly volatile environments, the utility of exploration may decrease rather than increase. This is because the predictive power of information gained through exploration–the value of exploration for making future decisions–will decay quickly. In this view, there should be a U-shape relationship between level of exploration and the volatility of the environment, rather than a linear one. We examined how subjects transitioned between exploration and exploitation in a restless three-armed bandit task with varying volatility levels. The results suggested that exploration was low at both low and extremely high volatility, despite the fact that performance was above chance in all cases. Fitting reinforcement learning models revealed low decision noise at high volatility, thus decision-making matched more closely with values of choices inferred from reward history. Together, these results demonstrated for the first time that exploration was a nonlinear function of volatility, where some intermediate volatility level maximize exploration.

2-094. Computational principles of motor adaptation

Barbara Feulner[1,2] B.FEULNER18@IMPERIAL.AC.UK
Matthew Perich[3] M@TTPERI.CH
Lee Miller[4] LM@NORTHWESTERN.EDU
Raeed Chowdhury[5] RAEED.CHOWDHURY@PITT.EDU
Juan Gallego[1] JUAN-ALVARO.GALLEGO@IMPERIAL.AC.UK
Claudia Clopath[1] C.CLOPATH@IMPERIAL.AC.UK

[1]Imperial College London
[2]Bioengineering
[3]Mount Sinai
[4]Northwestern University
[5]University of Pittsburgh

Animals, including humans, have a remarkable ability to adapt movements to changing external conditions. Despite advances in understanding motor adaptation at the behavioural level, disentangling the underlying neural mechanisms remains challenging. Dorsal premotor (PMd) and primary motor (M1) cortex are the main cortical areas involved in movement planning and execution, and neural activity changes within these two areas are believed to underlie motor adaptation. However, it remains unknown whether changes in neural activity emerge due to synaptic changes within the motor cortices or through altered inputs to them. As there are currently no experimental methods to directly measure synaptic strength in vivo, answering this question must rely on indirect evidence from neural recordings, obtained through correlation or more sophisticated inference analyses. Yet, the extent to which synaptic changes can be detected with these measures remains unclear. We used a new approach to investigate the role of synaptic changes in PMd and M1 for motor adaptation. Instead of trying to infer weight changes from experimental data, lacking ground truth knowledge, we implemented a recurrent neural network architecture that solved motor adaptation tasks through changing its connectivity. This allowed us to 1) measure task-related connectivity changes, 2) compare ground truth weight changes to the results of common inference methods and 3) compare resulting activity changes to experimental data. Using this approach, we found that two commonly studied motor adaptation paradigms – force-field adaptation and visuo-motor rotation – require only small changes in network activity and underlying connectivity to counteract the perturbations. Furthermore, matching experimental observations, the correlation structure among network units was highly preserved, which is usually interpreted as stable synaptic connectivity. Thus, experimentally observed activity changes in PMd and M1 during motor adaptation could potentially arise from small, but effective, synaptic weight changes in PMd and M1.

2-095. Observation to execution cross-decoding of hand grips from ventral premotor cortex spiking activity

Eleftheria Tzamali[1] ETZAMALI@GMAIL.COM
Vassilis Papadourakis[1] VPAPADOURAKIS@GMAIL.COM
Vassilis Raos[2,3] VASRAOS@GMAIL.COM

[1]University of Crete Medical School
[2]University of Crete
[3]Medical School

Our goal was to (i) explore whether hand grips can be decoded from spiking activity of the ventral premotor cortex

(PMv) and, most critically, (ii) investigate the reliability and efficiency of cross-decoding i.e. whether a model trained with the action observation-elicited activity can effectively and consistently decode hand configurations using activity during action execution. Mirror neuron (MirN) activity was recorded from the ventral premotor cortex (n=122) of macaque monkeys who either reached and grasped objects using distinct grips (n=4), or observed the same movements executed by the experimenter. Support vector machines (SVM) were used for decoding. Using the entire PMv population, the algorithm correctly predicts all grips, both in observation and execution conditions. When k-subset of neurons is randomly selected, 50 (30) neurons are sufficient to achieve 95% mean accuracy in observation (execution) across 100 runs. The size of the population required to reach high accuracy levels was also estimated using a greedy-selection procedure: the unit with the best performance was initially selected and at each subsequent step a unit was added to those of the previous step so that the performance of the resulting population was the highest. Following this procedure five and two neurons were sufficient to achieve 100% accuracy during observation and execution, respectively. The mean accuracy of cross-decoding using either the entire population or randomly selected ensembles was around 25%, close to random. However, the application of the greedy-selection procedure resulted in a cross-decoding accuracy of 97% using only 6 neurons. This is the first study to report high cross-decoding accuracies of hand configurations from spiking cortical activity. This finding renders PMv, and MirNs in particular, promising for the development of accurate decoders that can control distinctive hand shapes in applications where activity and behavior cannot be co-recorded as in the case of paralysis.

2-096. C. elegans nervous system model reveals a XOR-gate mechanism mediating oxygen sensation

Jimin Kim[1,2] JK55@U.WASHINGTON.EDU
Eli Shlizerman[1] SHLIZEE@U.WASHINGTON.EDU

[1]University of Washington
[2]Electrical and Computer Engineering

The nematode Caenorhabditis elegans (C. elegans) performs aerotaxis in which locomotion is modulated according to the composition of air in the environment. We extend a recently developed whole nervous system neuro-mechanical model of C. elegans to investigate aerotaxis mediated by oxygen sensing neurons AQR and PQR. Our method implements stimulation of any group of neurons (injection of in silico currents) according to the location of the body. We use the capability to construct spatial regions mimicking environmental oxygen presence, e.g., gradients, which indicate stimulation amplitudes of neurons distributed throughout the space. We show that the model can reproduce turning behavior under spatial regions of AQR and PQR stimuli as observed in in-vivo. Through systemic computational ablation, we find that interneuron PVPL plays an essential role in facilitating the turn for these sensors. When AQR and PQR are stimulated together, the turn behavior is inhibited and the worm slowly comes to stop, indicating that AQR and PQR form XOR gate-like mechanism to control the turn. By utilizing both systematic ablation and analyzing activity of simulated neural activities, we show that the XOR gate is implemented by the activities of PVPL and RIGL inter-neurons.

2-097. Neural circuits for converting sensory information into a motor plan

Vahid Esmaeili[1,2] VAHID.ESMAEILI@EPFL.CH
Keita Tamura[1] KEITA.TAMURA@EPFL.CH
Samuel Muscinelli[1] SAMUELEMUSCINELLI@GMAIL.COM
Alireza Modirshanechi[1] ALIREZA.MODIRSHANECHI@EPFL.CH
Marta Boscaglia[1] BOSCAGLIAMARTA@GMAIL.COM
Ashley Lee[1] ASHLEYMINLEE@GMAIL.COM
Anastasiia Oryshchuk[1] ANASTASIIA.ORYSHCHUK@EPFL.CH
Georgios Foustoukos[1] GEORGIOS.FOUSTOUKOS@EPFL.CH
Yanqi Liu[1] YANQI.LIU@EPFL.CH
Sylvain Crochet[1] SYLVAIN.CROCHET@EPFL.CH
Wulfram Gerstner[1] WULFRAM.GERSTNER@EPFL.CH
Carl Petersen[1] CARL.PETERSEN@EPFL.CH

[1]Ecole polytechnique federale de Lausanne
[2]SV-BMI-LSENS

Interpreting incoming sensory information in a context- and learning-dependent manner to act appropriately is a mental operation we perform continuously in daily life. Purposeful behavior often requires planning of motor actions in advance. The neural circuits processing sensory information or planned actions have often been studied

in isolation, in different behavioral tasks and by examining a few regions of interest. Moreover, the circuit principles that convert a sensory signal to a motor plan have not been elucidated. Here, we performed a comprehensive investigation of cortical neuronal circuits contributing to sensorimotor transformation in a single behavioral task and delineated important computations across learning. We trained mice to respond to a brief whisker stimulus by licking after a delay. The delay period temporally separated sensory processing and motor planning from action execution, allowing us to selectively study the learning signatures of these different mechanisms. We combined high-density extracellular recordings before and after task learning, spatiotemporally specific optogenetic inactivation, high-speed video filming of behavior and extensive computational quantifications to unravel the causal contribution of sensory, motor and higher-order cortical areas to different aspects of sensorimotor transformation. In particular, we found that sensory processing, more localized to whisker sensorimotor regions in Novice mice, distributed to broader brain areas upon learning. Motor planning engaged mostly the secondary tongue-jaw motor region (also referred to as anterolateral motor area, ALM), showing preparatory activity during the delay period, while primary tongue-jaw region was more active during the execution, and showed inhibitory response during the delay, when suppression of immediate lick was required. The secondary whisker motor cortex (wM2) showed the earliest choice related activity and learning-enhanced response to the whisker stimulus. These results suggest a specific cortical circuit with wM2 acquiring a pivotal role in transforming whisker information into preparatory activity for goal-directed motor planning.

2-098. Circuit motifs for vector navigation and goal-directed action: insights from Drosophila connectomics

Brad Hulse[1] BRADKHULSE@GMAIL.COM
Daniel Turner-Evans[1] TURNEREVANSD@JANELIA.HHMI.ORG
Hannah Haberkern[1] HABERKERNH@JANELIA.HHMI.ORG
Romain Franconville[1] FRANCONVILLER@JANELIA.HHMI.ORG
Shin-ya Takemura[1] TAKEMURAS@JANELIA.HHMI.ORG
Tanya Wolff[1] WOLFFT@JANELIA.HHMI.ORG
Marcella Noorman[1,2] NOORMANM@JANELIA.HHMI.ORG
Marisa Dreher[1] DREHERM@JANELIA.HHMI.ORG
Chuntao Dan[1] DANC@JANELIA.HHMI.ORG
Ruchi Parekh[1] PAREKHR@JANELIA.HHMI.ORG
Ann Hermundstad[1] HERMUNDSTADA@JANELIA.HHMI.ORG
Gerald Rubin[1] RUBING@JANELIA.HHMI.ORG
Vivek Jayaraman[1] VIVEK@JANELIA.HHMI.ORG

[1]HHMI Janelia Research Campus
[2]Computation and Theory

Despite a growing catalog of neurons that form internal representations suggestive of navigational computations, how these representations are generated and used to guide behavior is largely unknown. The navigation center of insects resides in a highly recurrent, evolutionarily conserved region known as the central complex (CX). The CX houses a ring attractor network that computes the fly's head direction (HD) using both self-motion and allothetic sensory cues [1-3]. Within this network, whose topography matches its function [4-5], HD cells are arranged into columns according to their preferred firing direction, and population activity is organized as a traveling activity 'bump' whose columnar position encodes the fly's HD. Here we analyze a recently released [6], complete connectome of the Drosophila CX and identify circuit motifs downstream of the fly's compass network whose structure strongly suggests the implementation of vector-based computations for navigation. We used these motifs to build a conceptual model that describes how cosine-formatted activity bumps could form a four-vector basis set for navigational computations (Figure 1). The motifs suggest that the amplitudes of the basis vectors could be independently gain-modulated [7] by self-motion signals to compute behaviorally relevant allocentric vectors. We highlight the utility of this motif in the context of path integration, describing how the four-vector basis set could compute an insect's translational velocity vector (i.e. its heading), which need not be the same as its HD. Finally, CX output neurons appear anatomically configured to compare the fly's current HD to that of an allocentric goal direction to generate egocentric motor commands. These results establish general circuit motifs for vector computation in the CX, with specific implementations that likely depend on cell type, species, and navigational need. More broadly, they highlight how the fly connectome can be used to generate functional hypotheses about broadly relevant navigational circuit computations.

2-099. Unsupervised learning of a dictionary of neural impulse responses from spiking data

Bahareh Tolooshams[1,2] BTOLOOSHAMS@SEAS.HARVARD.EDU
Hao Wu[1,3] HWU01@G.HARVARD.EDU
Paul Masset[1,4] PAUL_MASSET@FAS.HARVARD.EDU
Venkatesh Murthy[1] VNMURTHY@FAS.HARVARD.EDU
Demba Ba[1,2] DEMBA@SEAS.HARVARD.EDU

[1]Harvard University
[2]School of Engineering and Applied Sciences
[3]Molecular and Cellular Biology
[4]Department of Molecular and Cellular Biology

The majority of approaches for studying neuronal activity use an experimental setup where a stimulus is repeatedly applied over a series of trials with time-locked and non-overlapping events. Then, the spike trains are averaged over trials and smoothed out. These approaches fail in naturalistic environments and experiments in which the stimulus comprises discrete events occurring at random times, which may elicit overlapping responses. To analyze neuronal activity patterns in such experiments, we utilize a model of the spiking rate of a neuron as the convolution of an unknown impulse response and a sparse code, representing the time when the stimulus elicits an activity pattern in the neuron's response, and the response's amplitude. We fit the model to single-unit spiking data by solving a Poisson dictionary learning problem that lets us estimate a neuron's impulse response, and the amplitude of the response to each stimulus, directly from the spiking data. To solve the problem, we construct an autoencoder. We used neural spiking data acquired from piriform cortex in response to odor pulses to estimate the impulse responses (dictionary) of ~200 neurons along with the strength of the response associated with each pulse for each neuron (sparse code). The Kolmogorov-Smirnov (KS) test shows that the model fits the data well. Our analysis shows that, at the level of a single neuron, the odor pulses evoke different responses, likely reflecting differences in alignments to the breathing phase. In addition, we found that neurons from the population cluster according to either the estimated impulse responses or the stimulus responses, suggesting the presence of distinct neural populations in piriform cortex that could have distinct roles in the processing of information in olfactory search.

2-100. Towards replicating the mouse visual cortex in Intel's neuromorphic hardware

Srijanie Dey[1,2] SRIJANIE.DEY@WSU.EDU
Alexander Dimitrov[1] ALEX.DIMITROV@WSU.EDU

[1]WASHINGTON STATE UNIVERSITY
[2]Mathematics

Neuromorphic chips, as the name suggests - 'like the brain', can mimic the brain function in a truer sense as their design is inspired from the brain. Inspired by its architecture, we work on developing a principled approach towards obtaining simulations of biologically realistic brain networks on a novel neuromorphic hardware platform. The computer chips today are limited in this respect because of the way they have been built historically and the way they process data leading towards more energy and resource consumption. Neuromorphic chips on the other hand claim to be more efficient. We try to validate this assertion based on various brain networks derived from the primary visual cortex of the mouse. The work comes with its challenges as simulations built on the conventional chips cannot be trivially mapped to the neuromorphic platform. Since the neuromorphic chip is built based on the brain, the architecture differs remarkably from the conventional hardware and presents us with unique mapping challenges. For accessing the conventional hardware, we work with a software package called the Brain Modeling Toolkit (BMTK) developed by the Allen Institute of Brain Science (AIBS), while for our neuromorphic platform we use Intel Corporation's new experimental neuromorphic chip - 'Loihi'. We work with various measures to validate results between the two platforms. The success of this work will lay the foundation for replicating, and subsequently extending the brain's functionalities on hardware/machine/computer.

2-101. Motor cortical neural dynamics are finely spatially intermingled

Daniel O'Shea[1] DJOSHEA@STANFORD.EDU
Eric M Trautmann[2] EMT2177@COLUMBIA.EDU
Krishna Shenoy[3] SHENOY@STANFORD.EDU
Xulu Sun[1] XULU2@STANFORD.EDU
Saurabh Vyas[2] SMVYAS@STANFORD.EDU

[1]Stanford University
[2]Columbia University
[3]Boston University

Recent advances in silicon electrode technology enable dense, simultaneous sampling of many neurons. As significant computational and theoretical work has begun to understand the computations performed via neural population dynamics, these tools open experimental avenues to explore how these dynamics are organized within neural circuits. Using NeuroPixel probes for the first time in rhesus macaque primary motor and dorsal premotor cortices, we collected a dataset comprising 6,990 neurons (36 sessions, 2 monkeys). We analyzed neural responses during a reaching task to investigate the spatial, laminar, and synaptic organization of neural populations engaged in motor control. We first revisited a long-held view that motor cortex exhibits columnar architecture organized along shared preferred movement directions. A columnar organization predicts that spatially proximal neuron pairs should share preferred directions. Instead, we found uniformly distributed directional tuning independent of distance. More generally, spatial proximity was not predictive of similarity in neural responses assessed via PSTH correlation nor angles between GPFA loading vectors. A coarse laminar division of neurons into superficial and deep cells revealed increased similarity among superficial neurons, driven nearly entirely by similarity in the condition-independent component of neurons' movement responses and shorter latency to respond to task cues. Cross-correlation analysis revealed 5,639 / 843,033 (0.67%) neuron pairs with sharp, low-latency peaks after jitter correction, which removes slow-timescale structure. These neuron pairs exhibited more aligned preferred directions and more correlated PSTHs, reflecting increased similarity in both condition-independent and dependent response components. Collectively, these analyses depict highly heterogeneous, finely intermingled neuronal responses throughout primary and premotor cortex. Laminar structure is evident in the dominant, condition-independent response component, reflecting targeted innervation by corticocortical and thalamocortical afferents which initiate and guide movement dynamics. In contrast, condition-specific response dynamics are densely, spatially intermingled and organized along a network of spatially interwoven, synaptic interactions.

2-102. A nonlinear geometric dynamical modeling framework for neural population activity

Han-Lin Hsieh[1,2] HANLINHS@USC.EDU
Hamidreza Abbaspourazad[1] HABBASPO@USC.EDU
Bijan Pesaran[3] BIJAN@NYU.EDU
Maryam Shanechi[1] SHANECHI@USC.EDU

[1]University of Southern California
[2]Electrical and Computer Engineering
[3]New York University

Neural population dynamics in the motor cortex that underlie movements may evolve on low-dimensional yet nonlinear geometric manifolds. To date, developing machine learning methods that can explicitly identify and model nonlinear geometric manifolds underlying neuronal population dynamics remain elusive. Here, we develop a novel nonlinear geometric dynamical modeling framework for neural population activity that achieves this goal. A pervasive feature of motor cortical dynamics across studies is rotational patterns found with various dimensionality reduction methods. We hypothesize that neural population activity may evolve on nonlinear manifolds that contain holes, which may explain why linear projections onto hyperplanes naturally exhibit rotations in time. We thus developed a novel framework that can identify multi-dimensional (multi-D) nonlinear manifolds with holes based on neural data and then learn a dynamic model on top of these manifolds. We used our framework to study motor cortical population dynamics in two non-human primates (NHP) performing a 3D naturalistic reach-and-grasp task. We first tested our hypothesis and found that the manifold underlying neural dynamics was indeed nonlinear and contained a hole, yet was multi-D. Interestingly, the trajectory of neural population activity both along the hole and along other dimensions on the manifold were necessary to explain (decode) the movement. Finally, we trained an artificial recurrent neural network (RNN) to perform the same 3D reach-and-grasp task. We found that the RNN neurons also traversed a similar nonlinear manifold as the one found from motor cortical recordings. Taken together, we developed a new general tool for nonlinear dynamic modeling of neural population activity with the unique ability to identify the underlying geometry, and to build the dynamic model on top of it. By doing so, this framework enables nonlinear yet low-dimensional and interpretable modeling of neural population activity.

2-103. Motor cortex produces an output for behavior by using general activity transitions, not one pattern

Vivek Athalye[1,2] VA2371@COLUMBIA.EDU
Preeya Khanna[3,4] PREEYA.KHANNA@UCSF.EDU
Amy Orsborn[5,6] AORSBORN@UW.EDU
Suraj Gowda[7] SURAJGOWDA33@GMAIL.COM
Rui Costa[1] RC3031@COLUMBIA.EDU
Jose Carmena[7] JCARMENA@BERKELEY.EDU

[1]Zuckerman Mind Brain Behavior Institute, Columbia University
[2]Neuroscience
[3]University of California, San Francisco
[4]Neurology
[5]University of Washington
[6]Electrical & Computer Engineering, Bioengineering
[7]University of California, Berkeley

The temporal order of motor cortex output is important for controlling movement. Do cortical populations controlling movement produce an output using one specific neural population activity pattern, irrespective of the temporal order required for movement? This view is supported by studies explaining cortical activity as representation of behavioral parameters [1]. However, this is not necessary since many activity patterns can generate the same output [2]. Alternatively, network connectivity could cause activity patterns to have preferred transitions in time (dynamics), so that activity producing an output also encodes distinct previous and future activity [3]. If a complete motor cortex population and its output that controls behavior could be measured, we could test these alternatives. Accordingly, we used a brain-machine interface (BMI) to define the motor cortex neurons controlling movement, the decoder which linearly maps binned spike counts to two-dimensional motor cortex output, and the neuroprosthetic which assigns motor cortex output to the velocity of a two-dimensional cursor. We trained two rhesus macaques to perform straight and curved movements requiring different temporal orders of velocity outputs. Critically, we found that activity generating a specific output differed across movements. A linear model of activity dynamics predicted the distinct activity for the same output. Crucially, dynamics predicted appropriately distinct activity even for unseen outputs and movements, illustrating generality of dynamics. Further, distinct activity for the same output preferentially transitioned to future activity to produce distinct outputs appropriate for each movement. These results suggest that the brain produces ordered outputs not by re-using one pattern for each output, but by controlling activity to use preferred transitions. In support, we built a feedback control model of activity controlling the BMI, and found that experimentally-observed dynamics can help BMI control by reducing the required input to motor cortex and result in predictably distinct output-specific activity across movements.

2-104. Distributed representations in primate DLPFC with strategy use in a self-ordered working memory task

Feng-Kuei Chiang[1,2] FENG-KUEI.CHIANG@MSSM.EDU
Joni Wallis[3] WALLIS@BERKELEY.EDU
Erin Rich[1,4] ERIN.RICH@MSSM.EDU

[1]Icahn School of Medicine at Mount Sinai
[2]Nash Department of Neuroscience
[3]University of California, Berkeley
[4]Neuroscience

The limited capacity of working memory (WM) can be improved by using cognitive strategies. We previously showed that self-generated sequencing strategies reduced incorrect responses when monkeys performed a spatial target selection task. At the same time, the spatial tuning observed in dorsolateral prefrontal cortex (DLPFC) neurons decreased with more stereotyped sequencing [1]. Here, we assessed how this task is represented in neural ensembles, under these different self-generated behaviors. We re-analyzed data from two monkeys performing the self-ordered WM task with six identical visual targets. The task required subjects to saccade to each target, one at a time in any order, returning their eyes to the center after each selection. Reward was delivered only once for each target, and repeat visits were not rewarded. Therefore, monkeys had to use WM to track which targets had been visited and prepare for the next target selection. Sequencing strategies were identified as common selection patterns in blocks of trials where the target configuration was held constant. We found that target location and saccade number could be decoded from ensembles of DLPFC neurons. Interestingly, decoding performance was similar or better when monkeys used more stereotyped selection strategies, seemingly contradicting single unit results that found less spatial tuning. To investigate this further, we used multiple approaches to demonstrate that the informative contribution of single units to the ensemble decreased, and the optimal ensemble sizes for decoding increased under the self-generated sequencing strategy, suggestive of a more distributed but less efficient neural code. Lastly, we found that the informative dimensionality of neural ensembles was positively correlated with WM loads, and negatively correlated with stereotyped strategy. In sum, these data indicate that sequencing strategies change the local distribution of WM representations in DLPFC, with more stereotyped behaviors resulting from a more distributed neural code.

2-105. Hippocampal neurons use weighted combination of sensory inputs to encode spatial location

Andrey Sobolev[1,2] SOBOLEV.ANDREY@GMAIL.COM
Kay Thurley[3,4] THURLEY@BIOLOGIE.UNI-MUENCHEN.DE
Dustin Fetterhoff[1] FETTERHOFF@BIOLOGIE.UNI-MUENCHEN.DE
Christian Leibold[1] LEIBOLD@BIO.LMU.DE
Anton Sirota[1] SIROTA@BIOLOGIE.UNI-MUENCHEN.DE

[1]Ludwig Maximilian University of Munich
[2]Biology II
[3]LMU Munich
[4]Department Biology II

Hippocampal place cells show both stable representation of the environment as well as flexibility to changes of the features of the environment. Balance between stability and flexibility of spatial representations relies on the concurrent update of internal spatial representation by external sensory inputs and self-motion cues, however, it is still not known how exactly these inputs interact to build a stable representation of space or how their mismatch results in changes of the representation. Here we use the Virtual Reality system for freely-moving rodents that allows to investigate the effect of visual- and positional- manipulation on the hippocampal space code while keeping natural behaving conditions. Specifically, we investigate the effects of a conflict between visual-cue-defined and physical-boundary-defined reference frames on the hippocampal representation of space. Despite previous suggestions that place cells responses to visual and self-motion cues are organized in distinct categories (Fattahi et al., 2018, Chen et al., 2013), we provide experimental evidence that most of the place cells are involved in representing both reference frames using a weighted combination of sensory inputs. In line with the studies showing dominance of the more reliable sensory modality (Jeffery O'Keefe, 1999; Gothard et al., 2001), our data is consistent with CA1 cells implementing a maximum likelihood estimation given the idiothetic and allocentric inputs with weights inversely proportional to the availability of the input, as proposed for other sensory systems (Jeffery et al., 2016). This mechanism of weighted sensory integration, consistent with recent dynamic loop models of the hippocampal-entorhinal network (Li, Sheynikhovich et al., 2020), can contribute to the physiological explanation of Bayesian inference and optimal combination of spatial cues for localization (Cheng et al., 2007).

2-106. Dorsomedial frontal cortex participates in both evidence accumulation and history-based updating

Thomas Luo[1,2] THOMAS.ZHIHAO.LUO@GMAIL.COM
Tim Hanks[3] THANKS@UCDAVIS.EDU
Diksha Gupta[1,2] DIKSHAG@PRINCETON.EDU
Carlos Brody[2] BRODY@PRINCETON.EDU
Adrian Bondy[1] ABONDY@PRINCETON.EDU

[1]Princeton University
[2]Princeton Neuroscience Institute
[3]University of California, Davis

When making a decision based on noisy sensory stimuli, integration over time mitigates the impact of noise and is therefore a core component of many types of decisions. In trial-based perceptual tasks, the decision process is learned and continually updated through the evaluation of actions in past trials and their outcomes. The neural mechanisms underlying evidence accumulation and the updating of the decision process by trial history remain poorly understood. It is unknown whether the same brain region that mediates evidence integration in expert observers is also involved in updating the process of choice formation. To identify a candidate brain region, we first examined whether changes in history-dependent updating can be detected after the pharmacological inactivation of five separate sensorimotor brain regions in rats accumulating noisy auditory evidence. Only the perturbation of dorsomedial frontal cortex (dmFC) resulted in both an impairment in behavioral sensitivity to the sensory stimulus and also a change in trial history-dependent updating. Chronic lesion of dmFC similarly changed both behavioral sensitivity and history-dependent updating. Three separate experiments indicate that dmFC is

causally involved in the gradual accumulation of evidence. Recordings from Neuropixels probes indicate that the encoding of accumulated sensory evidence is more robust and emerges more rapidly in neural populations in dmFC than in four other anterior brain regions. Optogenetic inactivation of dmFC specifically during trial epochs when evidence was accumulated impaired behavioral sensitivity. Lastly, either optogenetic or chronic inactivation of dmFC resulted in an impaired representation of sensory evidence in anterior brain regions interconnected with dmFC. These results indicate the rat dorsomedial frontal cortex is a shared node in the circuits that mediate the gradual accumulation of evidence and history-dependent updating and suggests that activity in dmFC encodes action values that are used for both choice formation and modification of subsequent decisions.

2-107. Clusters structure spontaneous activity in presynaptic networks but not randomly sampled neurons

Trang-Anh Nghiem[1,2] TRANGANH.NGHIEM@GMAIL.COM
Stuart Trenholm[3] STUART.TRENHOLM@MCGILL.CA
Arjun Bharioke[4] BHARIOKE@GMAIL.COM
Adrian Wertz[5] ADRIAN.WERTZ@WEB.DE
Botond Roska[4] BOTOND.ROSKA@IOB.CH
Rava Azeredo da Silveira[6] RAVA@IOB.CH

[1]Ecole Normale Superieure
[2]Department of Physics
[3]McGill University
[4]Institute of Molecular and Clinical Ophthalmology Basel
[5]Friedrich Miescher Institute
[6]Ecole Normale Superieure, Institute of Molecular and Clinical Ophthalmology Basel

Our visual perception of the world around us emerges from rich interactions within networks of synaptically connected neurons, each neuron along the visual pathway only accesses information on stimuli encoded by its presynaptic network. However, it is unclear how neurons within a given presynaptic network collectively convey information, as a majority of previous studies on neural coding have randomly sampled subsets of neurons in a region. To address this issue, we monitored neuronal activity within the specific presynaptic neurons providing input to the same neuron in mouse primary visual cortex layer 4 by combining single-cell-initiated transsynaptic tracing with GCaMP6s-expressing rabies viruses and high-speed 3D 2-photon imaging. We find that while pairwise correlations can convey a range of statistical properties of spontaneous activity in randomly sampled neuronal populations, spontaneous activity among presynaptic networks is better captured by higher-order interactions. In particular, high-order interactions are found to be organised into clusters of co-active neurons within each presynaptic network. To investigate whether clusters defined from spontaneous activity are functionally relevant, we inspect the light-evoked responses from neurons within the same cluster. Spontaneous and evoked neural correlations appear aligned in randomly sampled neurons, but depart from each other in presynaptic networks. Our results highlight fundamental differences in neural correlation structure in presynaptic networks and randomly sampled cells, where presynaptic activity is described in terms of clusters defined by higher-order interactions among neurons.

2-108. Distributed synaptic plasticity accelerates learning in a multi-layer network of an electric fish

Salomon Z Muller[1] SZM2106@COLUMBIA.EDU
Nathaniel B Sawtell[1] NS2635@COLUMBIA.EDU
Larry Abbott[2] LFA2103@COLUMBIA.EDU

[1]Zuckerman Mind Brain Behavior Institute, Columbia University
[2]Columbia University

Understanding how learning arises from synaptic plasticity that is widely distributed in feedforward and recurrent circuits is a central challenge in neuroscience. Here we examine the specific contributions of multiple sites of synaptic plasticity to a well-described form of learning in the electrosensory lobe (ELL) of weakly electric mormyrid fish. Learning in the ELL enhances the processing of behaviorally relevant electrosensory signals by cancelling responses to the fish's own electrical field. Though past work has shown that such cancellation is mediated by anti-Hebbian synaptic plasticity at synapses onto both intermediate and output layer neurons, the specific contributions of multiple sites of plasticity within the ELL is unknown. Using modeling and experiments we identify specific constraints that limit the speed of cancellation based on a single site of anti-Hebbian plasticity at the output stage. Increasing the learning rate at the output stage introduces fluctuations in neuronal responses that

would interfere with sensory processing and are much larger than those observed experimentally. In contrast, adding learning at middle layer neurons, that are themselves recurrently connected, accelerates learning by ~10-20 fold, a time frame consistent with that observed experimentally. These results have relevance to quantifying rate of learning in brain circuits, and to understanding the role of multi-site plasticity that is a common feature of many brain circuits.

2-109. The recurrent neural circuits of contextual effects in visual cortex

Drew Linsley DREW_LINSLEY@BROWN.EDU
Lakshmi Govindarajan LAKSHMI_GOVINDARAJAN@BROWN.EDU
Thomas Serre THOMAS_SERRE@BROWN.EDU

Brown University

Neurons in the visual cortex are sensitive to context: Responses to stimuli presented within their classical receptive fields (CRFs) are modulated by stimuli in their surrounding extra-classical receptive fields (eCRFs). However, the circuits underlying these contextual effects are not well understood, and little is known about how these circuits drive perception during everyday vision. We tackle these questions by approximating circuit-level eCRF models with a differentiable discrete-time recurrent neural network that is trainable with gradient-descent. After optimizing model synaptic connectivity and dynamics for object contour detection in natural images, the neural-circuit model rivals human observers on the task with far better sample efficiency than state-of-the-art computer vision approaches. Notably, the model also exhibits CRF and eCRF phenomena typically associated with primate vision: (i) winner-takes-all selectivity, (ii) the orientation-tilt illusion as well as (iii) tuned and (iv) feature-specific surround suppression. The model's ability to accurately detect object contours also critically depends on these effects, and these contextual effects are not found in ablated versions of the model. Next, we sought to characterize the synaptic connectivity resulting from optimizing the model for contour detection. We developed a novel method for inferring in-silico functional connectivity between neurons, which involves stimulating a neural population and then inferring every other neural population's compensatory response to correct the stimulation's effect. This yields a map depicting the selected population's inhibitory and excititory eCRFs. The model's map recapitulates the pattern described in recent in-vivo optogenetics (Chettih & Harvey, 2019), with separate near-excitatory vs. far-inhibitory eCRFs. Moreover, by adjusting the stimulation strength, we find that the neural-circuit model controls contextual effects by changing the relative spatial extent of these opposing eCRFs. Overall, our work makes testable predictions about the neural mechanisms responsible for contextual integration and illustrates their importance for accurate and efficient visual perception.

2-110. Individual variability of neural mechanisms underlying flexible decision-making

Marino Pagan[1,2] MPAGAN@PRINCETON.EDU
Vincent Tang[2] VDTANG@PRINCETON.EDU
Mikio Aoi[3,4] MAOI@UCSD.EDU
David Sussillo[5] SUSSILLO@GOOGLE.COM
Valerio Mante[6] VALERIO@INI.UZH.CH
Jonathan Pillow[7] PILLOW@PRINCETON.EDU
Carlos Brody[2] BRODY@PRINCETON.EDU

[1]Princeton University
[2]Princeton Neuroscience Institute
[3]University of California, San Diego
[4]Data Science & Neurobiology
[5]Google / Stanford
[6]University of Zurich & ETH Zurich
[7]Princeton Neuroscience Institute, Princeton University

Neural mechanisms underlying an identical behavior can differ among individuals. Although individual variability has been studied in the context of stereotyped behaviors and simple sensorimotor associations, little is known about the heterogeneity of neural mechanisms supporting complex cognitive behaviors, which involve hidden internal variables. Here we show that rats can be trained to perform a sophisticated task requiring context-dependent selection and integration of noisy stimulus features, allowing the study of variability across multiple trained subjects. Our task is adapted from a closely related primate study that found that frontal regions such as the Frontal Eye Fields (FEF) play an important role in this process (Mante et al. 2013). Here, rats are presented with a train of randomly-timed auditory pulses, where each pulse varies in location (either right or left) and tone

pitch (high or low). On each trial, rats are cued to report either the prevalent location of the pulses, or their prevalent pitch. Optogenetic perturbation of rat Frontal Orienting Fields (FOF; putatively homologous or analogous to primate FEF) impaired rats' performance, consistent with a causal role of this region. Single-unit recordings revealed a striking similarity between population dynamics in rat FOF and previously reported dynamics in primate FEF. Taking advantage of the pulsatile nature of our stimulus, we developed analyses to distinguish between two competing models of context-dependent evidence selection, one relying on fast feedforward dynamics, the other on slower recurrent dynamics. Neural data displayed a surprising heterogeneity across subjects, revealing that each subject implements neural mechanisms that lie on a well-identified position along a spectrum between the two models. Validating our approach, predicted behavioral signatures of the neural mechanisms were in remarkably high agreement with estimates derived from neural data. Our approach opens the door to the study of individual variability in neural computations underlying higher cognition.

2-111. A diversity of discounting horizons explains ramping diversity in dopaminergic neurons

Paul Masset[1,2] PAUL_MASSET@FAS.HARVARD.EDU
Athar N. Malik[1,3] ATHAR.MALIK@MGH.HARVARD.EDU
HyungGoo R. Kim[1] HKIM07@G.HARVARD.EDU
Pol Bech[1] POLBECHV@GMAIL.COM
Naoshige Uchida[1] UCHIDA@MCB.HARVARD.EDU

[1]Harvard University
[2]Department of Molecular and Cellular Biology
[3]Massachusetts General Hospital

Temporal difference reinforcement learning (TD RL) has been extremely successful at characterizing the activity of dopaminergic neurons in the midbrain as signaling a reward prediction error (RPE). However, a number of recent experimental results from recordings in a wider set of tasks and anatomical locations have challenged this interpretation and showed a greater diversity of responses that cannot be explained by the canonical RPE framework. From a biological point of view, it seems unlikely that all the dopaminergic neurons convey the same signal. Here, we propose that distinct dopaminergic neurons have distinct discount factors and therefore compute prediction errors at different temporal horizons. Such diversity of discount factors has been recently shown to be theoretically advantageous and this interpretation can explain the qualitatively different ramping behavior of single neurons in a behavioral task within the normative RPE framework. We recorded the activity of optogenetically identified dopaminergic neurons in mice experiencing a 1-D virtual reality track. Although their average activity is consistent with an RPE signal, single neurons exhibit a structured diversity of responses that cannot be explained by traditional RL models. We show that a single convex value function (driven by the global reward expectation) experienced by neurons with a diversity of discount factors explains the diversity of these responses within the RPE framework. Our work extends a recent proposal that dopaminergic neurons perform distributional RL to another orthogonal dimension, learning across temporal horizons.

2-112. Learned motor patterns replayed in human motor cortex during sleep

Daniel Rubin[1,2] DRUBIN4@MGH.HARVARD.EDU
Tommy Hosman[3] TOMMY_HOSMAN@BROWN.EDU
Jessica Keleman[4] JKELEMEN@MGH.HARVARD.EDU
Brian Coughlin[4] BFCOUGHLIN@PARTNERS.ORG
Eyal Kimchi[4] EKIMCHI@PARTNERS.ORG
John Simeral[3] JOHN_SIMERAL@BROWN.EDU
Leigh Hochberg[4,3] LHOCHBERG@MGH.HARVARD.EDU
Sydney Cash[4] SCASH@MGH.HARVARD.EDU

[1]Massachusetts General Hospital, Harvard Medical School
[2]Neurology
[3]Brown University
[4]Massachusetts General Hospital

Consolidation of learning is believed to involve offline replay of neural activity. Though amply demonstrated in rodents, it is less well documented in humans, particularly regarding motor learning. Previous work has demonstrated evidence of replay during rest in human motor cortex immediately following a motor task, but no studies have explored this phenomenon during slow-wave sleep. To determine whether neural replay of learned motor sequences occurs during sleep, we recorded from human motor cortex as a research participant performed a novel

motor task and subsequently slept overnight. A 36-year-old man with tetraplegia secondary to cervical spinal cord injury had two 96-channel intracortical microelectrode arrays placed chronically into left pre-central gyrus (PCG) as part of an ongoing brain-computer interface pilot clinical trial. Single- and multi-unit activity was recorded while he played a color/sound matching memory game. On each of 160 trials, the participant was cued to a sequence of four screen locations; his task was to quickly move a computer cursor to those locations in the same sequence. Seventy-five percent of trials were the target sequence; the other trials were randomly interspersed distractor sequences. On each trial, intended movements were decoded by a real-time steady-state Kalman filter that allowed the participant to control a neurally-driven cursor on the screen. Intracortical neural activity from PCG and surface EEG were subsequently recorded overnight as he slept. When decoded using the same steady-state Kalman filter parameters, intracortical neural signals recorded overnight replayed the target sequence from the memory game throughout sleep (confirmed by surface EEG) significantly greater than chance. Replay occurred more often during slow-wave (NREM3) than NREM1/2 or REM sleep and approximately 10% faster than initial task performance. These results demonstrate that the neural activity associated with a recently learned task is replayed during sleep in human motor cortex.

2-113. Visual cortex is hardly auditory

Celian Bimbard C.BIMBARD@UCL.AC.UK
Timothy Sit TIM.SIT.18@UCL.AC.UK
Anna Lebedeva ANNA.LEBEDEVA.17@UCL.AC.UK
Philip Coen PIPCOEN@GMAIL.COM
Kenneth Harris KENNETH.HARRIS@UCL.AC.UK
Matteo Carandini M.CARANDINI@UCL.AC.UK

University College London

Sensory cortices are increasingly thought to encode multisensory information. For instance, mouse primary visual cortex (V1) appears to be influenced by auditory inputs, which have been suggested to provide global inhibition or loudness- and frequency-specific information. However, sounds can also evoke unsolicited behavioral responses, and such behavioral responses are now known to elicit neural activity across the brain. Thus, multisensory studies may suffer from a potential confound between acoustic-related and movement-related activity. We measured the neuronal responses to natural movies and natural sounds in V1 using chronically implanted Neuropixels probes, thus recording the activity of hundreds of units. At the same time, we filmed the behavior of the mice. We replicated previous results showing that sounds could evoke neural activity in V1. Though sound-evoked responses were weak compared to video-evoked responses, we could significantly decode sound identity from population activity. However, we also observed that in response to the sounds the mice performed stereotypical, unsolicited movements. These movements induced activity in the V1 population activity. Removing it significantly reduced the responses to sounds. We conclude that a significant part of sound-evoked responses in V1 comes from indirect behavioral responses, and that a careful characterization of auditory inputs to sensory areas requires controlling sound-induced changes in behavioral state.

2-114. Spike-constrained neural control

Itzel Olivos Castillo[1] ICO1@RICE.EDU
Paul Schrater[2] SCHRATER@UMN.EDU
Xaq Pitkow[3] XAQ@RICE.EDU

[1]Rice University
[2]University of Minnessota
[3]Baylor College of Medicine, Rice University

We develop a version of stochastic control that accounts for computational costs in the brain. Motor control and reinforcement learning both appeal to the conceptual framework of accumulating evidence about the world and selecting actions based on the synthesized information to maximize total expected utility. Yet neither of these approaches consider the costs of performing computations. Conversely, past studies identified metabolically efficient ways of coding sensory information, but these studies are restricted to feedforward settings and static environments, and do not consider the consequences for closed-loop control. Here we combine concepts of efficient coding with control theory to analyze Linear Quadratic Gaussian (LQG) control, a well-understood mathematical example of optimal control that combines a Kalman filter for partially observed stochastic linear dynamics of a gaussian world state, with a linear regulator that minimizes integrated quadratic state costs and action costs. We implement the Kalman filter neurally using a dynamic Probabilistic Population Code, in which linear projections of spiking neural activity approximate the natural parameters of a Gaussian posterior over the world state. To this framework we add a cost on the total integrated number of spikes. In the simplest version, the precision of

the inference is directly proportional to the number of spikes. This creates a trade-off: an agent can obtain more utility overall by relinquishing some task performance, if doing so saves enough spikes. By solving this problem, we describe how the optimal spike rate varies with the properties of the system to be controlled, such as stability, process noise, and observation noise. We discuss differences between how the brain should efficiently allocate resources in a closed-loop setting compared to the conventional feedforward setting. Overall, this work provides a foundation for a new type of bounded rational behavior that could be used to explain suboptimal computations in the brain.

2-115. Learning function from structure in neuromorphic networks

Laura Suarez[1] LAURASUAREZ24@GMAIL.COM
Blake Richards[1] BLAKE.RICHARDS@MILA.QUEBEC
Guillaume Lajoie[2] G.LAJOIE@UMONTREAL.CA
Bratislav Misic[3] BRATISLAV.MISIC@MCGILL.CA

[1]McGill Unviersity - Mila
[2]Universite de Montreal, Mila
[3]McGill University - Montreal Neurological Institute

Recent imaging technologies allow the reconstruction of the human connectome, a network model of the macroscale wiring patterns of the brain. Descriptive analyses of the connectome have found evidence of non-random architectural features, characteristic of complex systems, that theoretically shape the computational capacity of the brain. Features include: high clustering and short-path length, specialized segregated communities, and heavy-tailed degree distributions. Still, how network organization supports information-processing remains unknown. Broadly, we address the structure-function relationship in the brain, but with a focus on computation. To do so, we combine connectomics and reservoir computing to investigate the link between macroscale connectivity, and the computational properties that emerge from network dynamics in the human connectome. Specifically, we construct artificial neural networks endowed with biologically realistic connection patterns derived from MRI. We train these connectome-informed reservoirs to perform a memory task. To evaluate how performance depends on network structure and dynamics, we parametrically drive the network to transition between stable, critical and chaotic states. Throughout, we assess computational performance of empirically-derived connectomes against two null network models. We find that: i) the underlying macroscale topology and mesoscale modular organization of the brain enhances computational performance in the context of critical dynamics; ii) the modular organization of the brain in functional systems constitutes a computationally relevant feature of the human connectome; iii) the connectome's network topology optimizes the trade-off between computational capacity and metabolic costs, and iv) it supports functional specialization by promoting regional heterogeneity of information content. Remarkably, we observe a prominent interaction between network structure and dynamics, such that the same underlying architecture can support a wide range of learning capacities across dynamical regimes. Besides the novelty of the methods, the main innovation lies on the shift from a phenomenological concept of function, towards a focus on how brain circuits support computational capacity.

2-116. An arousal gated visual circuit controls pursuit during Drosophila courtship

Tom Hindmarsh Sten[1,2] THINDMARSH@ROCKEFELLER.EDU
Rufei Li[1] RLI01@ROCKEFELLER.EDU
Adriane Otopalik[1] AOTOPALIK@ROCKEFELLER.EDU
Vanessa Ruta[1] RUTA@ROCKEFELLER.EDU

[1]The Rockefeller University
[2]Laboratory of Neurophysiology and Behavior

Long-lasting internal states pattern ongoing behavior by defining how the sensory world is translated to specific actions that subserve the needs of an animal. Yet how enduring internal states shape sensory processing or behavior has remained unclear. Here, we used the Drosophila courtship ritual as a powerful inroad to explore how a persistent internal state restructures sensorimotor processing to guide behavior. At the onset of courtship, a male transitions from being 'blind' and apathetic to the presence female to faithfully tracking and singing to her for many minutes to persuade her to copulate. Combining functional imaging, quantitative behavior, and computational modeling, we gain insight into how the salience of female visual cues is transformed by a male's internal arousal state to give rise to persistent courtship pursuit. We reveal that the intensity of a male's arousal is continuously encoded in the activity of male-specific P1 neurons, which remain persistently active during bouts of courtship. This P1-mediated arousal serves to dynamically modulate the gain of LC10a visual projection neurons, which are necessary for a male's faithful pursuit. During arousal, visual signals are propagated with near lossless efficiency from LC10a neurons to descending neurons, giving rise to a male's faithful pursuit of the female

target. Indeed, a simple network model based on the LC10a circuit accurately predicts the moment-to-moment tracking behavior of a freely courting male over hundreds of seconds, underscoring that this circuit is nearly fully sufficient to explain a male's ongoing behavior during courtship. Together, these results reveal how alterations in a male's internal arousal state can dynamically modulate the propagation of visual signals through a high-fidelity visuomotor circuit, transforming a female from an indifferent visual object to a target of desire.

2-117. Mixed representations in a visual-parietal-retrosplenial network for flexible navigation decisions

Shin Kira[1] SHIN_KIRA@HMS.HARVARD.EDU
Houman Safaai[2] HOUMAN_SAFAAI@HMS.HARVARD.EDU
Ari Morcos[1] ARIMORCOS@GMAIL.COM
Giuseppe Pica[3] GPICA89@GMAIL.COM
Stefano Panzeri[3] STEFANO.PANZERI@IIT.IT
Christopher Harvey[1] CHRISTOPHER_HARVEY@HMS.HARVARD.EDU

[1]Harvard Medical School
[2]Harvard Medical School, Istituto Italiano di Tecnologia
[3]Istituto Italiano di Tecnologia

The survival of many animals requires navigation by making decisions based on sensory inputs and past experience. These navigation decisions arise from the flexibility to respond to the same sensory input with distinct actions depending on recent memories. However, navigation decisions have often been studied in tasks that require a stereotyped action in response to a given sensory stimulus, precluding investigation into the flexibility of decision-making. Here, we identified cortical areas and neural activity patterns essential for combining short-term memory with visual signals to inform navigation decisions in mice. We trained mice to perform a delayed match-to-sample task in a virtual reality T-maze. As a mouse ran through the maze, it sequentially observed two cues separated by a short delay (1-2 s). The mouse was required to combine the short-term memory of the first cue and visual information of the second cue to choose an appropriate turn direction at the T-intersection. For a given visual stimulus in the second cue, the mouse responded flexibly with a different turn direction based on the memory of the first cue. We performed systematic screening of brain areas and activity patterns critical for flexible decision-making. An optogenetic inhibition screen identified V1, posterior parietal cortex (PPC), and retrosplenial cortex (RSC) as necessary for accurate task performance. Two-photon calcium imaging revealed that individual RSC neurons efficiently mixed memory and visual information. This mixing occurred as distinct activity patterns for each combination of visual cues and memories, a format useful for guiding navigation. The mixed representations appeared to govern accurate decision-making because they were prominent before correct choices but degenerated during errors. We propose a mechanism for flexible decision-making during navigation based on mixing memory and visual signals primarily in RSC neurons to guide choices, within a visual-parietal-retrosplenial network.

2-118. Real-time neural feedback of mesoscale cortical GCAMP6 signals for training mice

Pankaj K. Gupta[1,2] PANKAJGUPTA@ALUMNI.UBC.CA
Timothy Murphy[1] TIMOTHY.MURPHY@UBC.CA

[1]University of British Columbia
[2]Department of Psychiatry

Mice can learn to control specific neuronal ensembles using sensory (eg. auditory) cues (Clancy et al. 2014) or even artificial optogenetic stimulation (Prsa et al. 2017). In the present work, we measure mesoscale cortical activity with GCaMP6s and provide graded auditory feedback (within ~100 ms after GCaMP fluorescence) based on changes in dorsal-cortical activation within specified regions of interest (ROI)s with a specified rule. We define a compact, low-cost optical brain-machine-interface (BMI) capable of image acquisition, processing, and conducting closed-loop auditory feedback and rewards, using a Raspberry Pi (Fig. 1). The changes in fluorescence activity (ΔF/F) are calculated based on a running baseline (eg. 5 sec.). Two ROIs (R1, R2) on the dorsal cortical map were selected as targets. We started with a rule of 'R1-R2' (ΔF/F of R1 minus ΔF/F of R2) where the activity of R1 relative to R2 was mapped to frequency of the audio feedback (Fig. 1D) and if it were to cross a set threshold, a water drop reward is generated. To investigate learning in this context, water-deprived tetO-GCaMP6s mice (N=TODO) were trained for 30-minutes every day on the system for several days, with a task to increase audio frequency leading to reward. We found that mice could modulate activity in the rule-specific target ROIs to get an increasing number of rewards over days (Figure 2C). Analysis of the reward-triggered ΔF/F over time indicated

that mice progressively learned to activate the cortical ROI to a greater extent (Figure 2B, 2A). In conclusion, we developed an open-source system (to-be released) for closed-loop feedback that can be added to experimental scenarios for brain activity training and could be possibly effective in inducing neuroplasticity.

2-119. The spectrum of covariance matrices of randomly connected recurrent neuronal networks

Yu Hu[1,2] MAHY@UST.HK
Haim Sompolinsky[3] HAIM@FIZ.HUJI.AC.IL

[1]The Hong Kong University of Science and Technology
[2]Mathematics
[3]Hebrew University; Harvard University

A key question in theoretical neuroscience is the relation between connectivity structure and the collective dynamics of a network of neurons. Here we study the connectivity-dynamics relation as reflected in the distribution of eigenvalues of the covariance matrix of the dynamic fluctuations of the neuronal activities, which is closely related to the network's Principal Component Analysis (PCA) and the associated effective dimensionality. We consider the spontaneous fluctuations around a steady state in a randomly connected recurrent network of stochastic neurons. An exact analytical expression for the covariance eigenvalue distribution in the large network limit can be obtained using results from random matrices. The distribution has a finitely supported smooth bulk spectrum and exhibits an approximate power law tail for coupling matrices near the critical edge. We then generalized the results to include connectivity motifs and to E-I networks. To facilitate empirical applications of the theory, we studied the effects of temporal and spatial sampling common to many large-scale neural recordings. Our results suggest that the covariance spectrum is a robust feature of population dynamics in recurrent neural circuits and provides theoretical predictions for this spectrum in simple connectivity models that can be compared with experimental data.

2-120. Noisy gradient updates drive dimensionality compression in neural networks

Matthew Farrell[1,2] MSFARRELL@SEAS.HARVARD.EDU
Stefano Recanatesi[3] STEFANO.RECANATESI@GMAIL.COM
Madhu Advani[1] M.ADVANI81@GMAIL.COM
Guillaume Lajoie[4] LAJOIE@DMS.UMONTREAL.CA
Eric Shea-Brown[3] ETSB@UW.EDU

[1]Harvard University
[2]Applied Mathematics
[3]University of Washington
[4]University of Montreal / MILA

When recording from animal brains during task engagement, researchers frequently observe that the responses of neurons in higher-order, task-driven brain areas occupy subspaces with a dimension that matches the number of task-relevant variables (for a review, see [1]). Here we elucidate this phenomenon using deep and recurrent neural network models trained to classify high-dimensional input data. We observe that, when trained with stochastic gradient descent, these networks spontaneously form task-dimensional representations through training. We argue that this phenomenon depends critically on noisy updates during gradient learning. We use mathematical arguments for a linear network to show how noise drives dimensionality compression, and test our theory in simulations using a variety of network models and optimization algorithms. Our tests indicate that dimensionality compression in many cases does not occur when gradient updates are deterministic (such as when using gradient descent with batch size equal to the number of training points). On the other hand, adding noise either by using minibatches or by injecting noise directly does, in many cases, drive dimensionality compression. This suggests that the noise ubiquitous in biology plays a key role in the formation of low-dimensional neural representations.

[1] Peiran Gao, Eric Trautmann, Byron Yu, Gopal Santhanam, Stephen Ryu, Krishna Shenoy, and Surya Ganguli. A theory of multineuronal dimensionality, dynamics and measurement. bioRxiv, 2017.

3-001. Hippocampal sharp wave-ripples and the associated sequence replay in a network model of area CA3

Andras Ecker[1,2] ANDRAS.ECKER@EPFL.CH
Bence Bagi[3] B.BAGI20@IMPERIAL.AC.UK
Eszter Vertes[4] VERTES.ESZTER@GMAIL.COM
Orsolya Steinbach-Nemeth[5] ORSI.NEMETH@HOTMAIL.COM
Tamas F. Freund[6] FREUND@KOKI.HU
Attila I. Gulyas[6] GULYAS@KOKI.HU
Szabolcs Kali[6] KALI@KOKI.HU

[1]Ecole polytechnique federale de Lausanne
[2]School of Life Sciences
[3]Imperial College London
[4]University College London
[5]Bosh
[6]Hungarian Academy of Science

Hippocampal place cells have spatial receptive fields that are sequentially activated as a rodent explores a novel environment. Replay of these neuronal activity sequences during resting states is essential for memory formation. To understand the link between learning from theta sequences during awake exploration and the sharp wave-ripple (SWR) dynamics of the resting period, we built a simplified network model of area CA3. Activity sequences from simulated exploration of a novel linear track were stored in the recurrent excitatory synapses of the CA3 network by the experimentally measured symmetric spike-timing-dependent plasticity (STDP) rule. Simulations of the resting period equipped with this learned recurrent weight structure displayed spontaneously occurring bursts of activity featuring sequence replays in both forward and backward directions, consistent with experimental observations. Replays were always accompanied by transient ripple frequency oscillations generated in the interconnected network of parvalbumin-positive basket cells (PV+BCs). Using different modifications of the recurrent excitatory weight matrix, we show that not the overall weight distribution, but the fine wiring structure is the key aspect for SWR-like network dynamics in our model. Further manipulations revealed that it was the combination of the symmetric STDP rule and the intrinsic adaptation of pyramidal neurons that enabled bidirectional sequence replay in our network. These results from our unifying model provide a more causal understanding of the links between memory storage, recall and hippocampal oscillations.

3-002. 3D animal pose estimation with a hierarchical von Mises-Fisher-Gaussian model

Libby Zhang[1] EYZ@STANFORD.EDU
Tim Dunn[2] TIMOTHY.DUNN@DUKE.EDU
Jesse Marshall[3] JESSE.D.MARSHALL@GMAIL.COM
Bence Olveczky[3] OLVECZKY@FAS.HARVARD.EDU
Scott Linderman[1] SCOTT.LINDERMAN@STANFORD.EDU

[1]Stanford University
[2]Duke University
[3]Harvard University

Animal pose estimation (APE) from video data is an important step in many biological studies, but current methods struggle in complex environments where occlusions are common and training data is scarce. Recent work using deep neural networks has demonstrated improved accuracy, but these methods often do not incorporate prior distributions that could improve localization. We present GIMBAL: a hierarchical von Mises-Fisher-Gaussian model that improves upon deep networks' estimates by leveraging spatiotemporal constraints. The spatial constraints arise from the animal's skeleton and induce a curved manifold of keypoint configurations. The temporal constraints arise from the postural dynamics and govern how directions between keypoints change over time. Importantly, the conditional conjugacy of the model permits simple and efficient Bayesian inference algorithms. We assess the model on a multiview video dataset of a freely-behaving rodent with corresponding ground-truth motion capture data, and compare its performance to two existing APE methods. We show how GIMBAL extends existing techniques, and in doing so, offers more accurate estimates of keypoint positions.

3-003. How can shared concept cells encode associations?

Chiara Gastaldi[1,2] CHIARA.GASTALDI@EPFL.CH
Tilo Schwalger[3,4] TILO.SCHWALGER@BCCN-BERLIN.DE
Wulfram Gerstner[1] WULFRAM.GERSTNER@EPFL.CH

[1]Ecole polytechnique federale de Lausanne
[2]Brain Mind Institute
[3]TU Berlin
[4]Institute for Mathematics

Experimental evidence suggests that memories of concepts are stored in the Medium Temporal Lobe (MTL) (Quiroga et al. 2005, Ison et al. 2008). Each concept is represented by a sparse cell assembly (or memory engram) so that only a small fraction of neurons responds to each concept. Associations between different concepts are related to shared concept cells: assemblies representing two arbitrary concepts share less than 1% of neurons, whereas assemblies representing previously associated concepts share up to 4% of neurons (De Falco et al. 2016). Attractor neural networks are the most commonly used tool for modeling associative memory. However, most literature only considers independent memory engrams with relative high fraction of active neurons, while experimental studies in the MTL measured 1. a low fraction of active neurons of about γ = 0.2% (Ison et al. 2015) and 2. a fraction of shared neurons between associated engrams significantly above chance. This makes the introduction of correlation between low-activity memory patterns a key missing piece of the theory. Using mean-field approximation, we derive analytic equations for the network dynamics in the case of correlated patterns in order to answer the following questions: How do shared neurons encode associations? How many neurons can be shared? We demonstrate that correlations between memory engrams induced by shared concept cells modify the memory recall process and provide a robust way to encode associations between pairs of concepts. Moreover, we show that the free recall of a chain of concepts is possible in very sparse attractor networks only if engrams are correlated. We quantify the critical values of correlation 1) above which patterns are indistinguishable during retrieval and 2) below which free recall of a chain of concepts is not possible. Our model explains electrophysiological findings in MTL (Quiroga et al. 2005, Ison et al. 2008) and links to, as well as extends, the theory of free recall of lists of words (Naim et al. 2020).

3-004. A neural model of ultra-fast learning with natural language instructions

Reidar Riveland[1,2] REIDAR.RIVELAND@GMAIL.COM
Alex Pouget[1] ALEX.POUGET@GMAIL.COM

[1]University of Geneva
[2]Neuroscience

In simple laboratory experiments, when humans are provided with a set of verbal instructions they can learn to perform unseen tasks in very few trials, while non-linguistic animals may require weeks of supervised training on the same tasks. Attempts to account for humans' remarkable adaptability often center on two factors: 1) the ability of the prefrontal cortex (PFC) to create abstract task representations that are subsequently composed to match the demands of novel settings and 2) the compositional nature of language which aids our extraordinarily flexible cognitive capacities. These are fundamental phenomena in cognitive neuroscience, yet a computational theory of how verbal instructions can aid in rapid learning remains elusive. Here we study the role of language in fast adaptability by training recurrent network models to perform many cognitive tasks simultaneously. In our control network, information about task type is represented as a one-hot indexing vector. In models with linguistic inputs, we take advantage of recent advances in natural language processing (NLP) and embed instructions using transformer architectures that have been pre-trained on a variety of NLP objectives. We find that in task holdout tests, our best language based model learns unseen tasks in dramatically fewer training examples than the control model. Often, near maximal performance is attained in a zero-shot manner. Importantly, instruction embeddings that facilitate such performance require that transformers are pre-trained on sentence-level semantic reasoning. Training only on predicting the next word of a sentence, as is done in GPT for instance, does not aid in generalization. Lastly, we find that rapid transfer of performance in instructed models is supported by the structured geometry of recurrent network hidden activity prior to stimulus onset. Holdout instruction embeddings initialize hidden units in the appropriate region of activity space relative to training tasks, thereby facilitating rapid adaptation of performance.

3-005. Transient chaotic SNR amplification

Christian Keup[1,2] C.KEUP@FZ-JUELICH.DE
Tobias Kuhn[3] T.KUEHN@PHYSIK.RWTH-AACHEN.DE
David Dahmen[1] D.DAHMEN@FZ-JUELICH.DE
Moritz Helias[4] M.HELIAS@FZ-JUELICH.DE

[1]Research Centre Julich
[2]INM-6
[3]ENS/CNRS Paris & Research Centre Julich
[4]Research Centre Julich & RWTH Aachen

Strongly chaotic non-linear networks strongly separate inputs, but are believed to be useless for classification tasks because also irrelevant (noise) differences within any class are exacerbated, leading to bad generalization. We show this is actually not the case during the initial time period following input presentation: During this time, the representation is dominated by expansion, but not by mixing, and larger differences (between classes) expand faster than smaller differences (within classes). Therefore, the representation is disentangled by the dynamics, and when classifying the network state by linear readouts, the signal-to-noise ratio (SNR) actually increases, before it eventually deteriorates when mixing begins to dominate. We show that this is a general effect in high-dimensional non-linear chaotic systems, and demonstrate it in spiking, continuous rate, and LSTM networks. The transient SNR amplification is always fast (within 50 ms) for spiking networks, while its timescale in continuous valued networks depends on the distance to the edge of chaos. Moreover, this fast, noise-resilient transient disentanglement of representations is in line with empirical evidence: the olfactory bulb, for example, rapidly enhances the separability of sensory representations in a single recurrent layer, being the initial processing stage of a relatively flat hierarchy.

3-006. Context-dependent action-memory representations in prefrontal cortex

Ioana Calangiu[1,2] CALANGIU@INI.UZH.CH
Valerio Mante[1] VALERIO@INI.UZH.CH

[1]University of Zurich & ETH Zurich
[2]Institute for Neuroinformatics

Memories are representations of the past that help us make better decisions in the future. A brain structure that is critical to maintaining short-term memories and planning future behavior is the prefrontal cortex (PFC). We performed recordings from large populations of neurons in PFC of monkeys engaged in a variety of tasks requiring saccades to collect rewards. Surprisingly, PFC responses were most strongly modulated after each saccade. We revealed the properties of post-saccadic activity by requiring monkeys to maintain fixation not just before, but also after each instructed saccade. Post-saccadic activity does neither reflect the planning of the subsequent saccade, nor the current gaze location. Rather, its properties are consistent with those of an action-memory, i.e. a sustained representation of the metrics of the previously performed action. While action memories are maintained in PFC for varying durations after every saccade (instructed or free), its strength is contextually modulated—action memories are strongest after saccades expected to lead to a reward, as well as in a task requiring the monkey to adapt its behavior based on the collected rewards. Overall, the properties of the identified action memory make it a candidate for an ‘eligibility trace’, a signal required by many theories of learning that so far lacked a clear counterpart in biology.

3-007. Generalized neural decoders for transfer learning across participants and recording modalities

Steven Peterson[1,2] STEPETER@UW.EDU
Zoe Steine-Hanson[1] ZSTEINEH@CS.WASHINGTON.EDU
Nathan Davis[1] NDAVIS98@UW.EDU
Rajesh P. N. Rao[1] RAO@CS.WASHINGTON.EDU
Bingni Brunton[1] BBRUNTON@UW.EDU

[1]University of Washington
[2]Biology

Advances in neural decoding have enabled brain-computer interfaces to perform intricate, clinically-relevant tasks.

However, such decoders are frequently trained on specific participants, days, and recording sites, limiting their practical long-term usage. Therefore, a fundamental challenge is to develop neural decoders that can robustly train on pooled, multi-participant data and generalize to unseen participants. Here, we introduce a new decoder, HTNet, that fuses deep learning with neural signal processing insights. Specifically, HTNet augments a pre-existing convolutional neural network decoder with two innovations: (1) a Hilbert transform that computes spectral power at data-driven frequencies and (2) a layer that projects electrode-level data onto predefined brain regions. This projection step is critical for intracranial electrocorticography (ECoG), where electrode locations are not standardized and vary widely across participants. We trained HTNet to decode arm movements using pooled multi-participant ECoG data and tested performance on an unseen ECoG or scalp electroencephalography (EEG) participant; these pretrained models were also subsequently fine-tuned to each test participant. We show that HTNet significantly outperformed state-of-the-art decoder accuracies by 8.5% when tested on unseen ECoG participants and by 14.5% when tested on unseen participants that used a different recording modality (EEG). We then used transfer learning and fine-tuned these trained models to unseen participants, even when limited training data was available. Importantly, we fine-tuned trained HTNet decoders with as few as 50 ECoG or 20 EEG events and still achieved decoding performance approaching that of randomly-initialized decoders trained on hundreds of events. Furthermore, we demonstrate that HTNet extracts interpretable, physiologically-relevant features. By generalizing to new participants and recording modalities, robustly handling variable electrode placements, and fine-tuning with minimal data, HTNet is more applicable across a broader range of neural decoding applications than current state-of-the-art decoders.

3-008. Distinguishing theories of sensory processing using neuronal activity perturbations

Sarah Armstrong[1,2] SARAH.ARMSTRONG@SEH.OX.AC.UK
Adam Packer[1] ADAM.PACKER@DPAG.OX.AC.UK
Andrew Saxe[3] ANDREW.SAXE@PSY.OX.AC.UK

[1]University of Oxford
[2]Department of Experimental Psychology
[3]Department of Experimental Psychology, University of Oxford

A variety of theories have been posited to explain neural responses in early sensory cortices including efficient coding, predictive coding, and supralinear stabilized networks (SSNs). Often these theories predict similar evoked sensory responses, raising a key theoretical question: what experiments could distinguish between these alternatives, revealing the underlying algorithm and circuit connectivity? Here, we develop a theoretical framework to explore the utility of neuronal activity perturbations to reveal computational principles of sensory processing. Under the assumption of approximately regular tuning of neurons to stimulus features in sensory cortices, we derive explicit 'influence functions' that describe how the network response to a single-neuron perturbation depends on tuning coefficients for a range of cortical theories and network architectures. We show that, while many theories are not identifiable from sensory steady state responses alone, they make experimentally distinctive predictions when both perturbations and responses are available. In particular we identify a novel link between influence functions and feedforward sensory tuning that is a hallmark of efficient coding; and we find differential effects of perturbing 'state' and 'error' neurons which distinguish predictive coding networks. Finally, we show in simulation how our results, which are based on a linearized analysis, can inform nonlinear processing, using sparse coding and the SSN as examples. Our results show how a combination of perturbations, sensory responses, and theoretical considerations can help constrain long-standing debates about the function and operating regime of early sensory cortices.

3-009. Behavioral and neural substrates of learning attentional rules

Caroline Jahn CAROLINE.ISABELLE.JAHN@GMAIL.COM
Britney Morea BMOREA@PRINCETON.EDU
Timothy Buschman TBUSCHMAN@PRINCETON.EDU

Princeton Neuroscience Institute, Princeton University

We must constantly adapt the rules we use to guide our attention. To understand how the brain learns attentional rules, we designed a novel task that required monkeys to learn which color is the most rewarded at a given time (the current 'attentional rule'). However, just as in real life, the monkey was never explicitly told the rule. Instead, they had to learn it through trial and error by choosing a color, receiving feedback (amount of reward), and then updating their internal attentional rule. Then, after the monkeys reached a behavioral criterion, the rule changed. This change was not cued but could be inferred based on reward feedback. Behavioral modeling found monkeys used rewards to learn attentional rules. After the rule changed, animals adopted one of two strategies. If the change in the rule was small, reflected in a small reward prediction error, the animals continuously updated their rule. However, for large changes in the rule (large reward prediction errors), monkeys ‘reset’ their belief about the rule and re-learned the rule from scratch. To understand the neural correlates of learning new attentional rules, we recorded from the prefrontal and parietal cortex. Preliminary results suggest prefrontal cortex neurons encode the rule, developing stable rule representations early in learning and maintaining them throughout the block of trials. This representation was lost immediately after the rule switch, reflecting the reset seen in behavior. In contrast, parietal cortex encoded the rule only once the rule was well-known and behavioral accuracy was high. Together, our results provide insight into the behavioral and neural basis of learning new abstract rules.

3-010. Dual wavelength mesoscopic imaging reveals spatiotemporally heterogeneous coordination of cholinergic and neocortical activity

Andrew Moberly[1,2] ANDREW.MOBERLY@YALE.EDU
Sweyta Lohani[1] SWEYTA.LOHANI@YALE.EDU
Hadas Benisty[1] HADAS.BENESTI@YALE.EDU
Boris Landa[1] BORIS.LANDA@YALE.EDU
Michael Higley[3] MICHAEL.HIGLEY@YALE.EDU
Jessica Cardin[3] JESS.CARDIN@YALE.EDU

[1]Yale University
[2]Neuroscience
[3]Yale School of Medicine

During wakefulness, animals cycle through brain states that profoundly influence perception and behavior, such as arousal, locomotion, and attention. A classical view suggests these states are global (e.g., homogeneously influencing the entire neocortex) and tightly linked to the output of neuromodulatory systems. Cholinergic neurons, which send widespread projections throughout the neocortex, are thought to play a key role in state-dependent modulation of neural activity. However, recent anatomical work has highlighted spatial diversity in cholinergic axonal arborizations, suggesting that these modulatory signals may be spatially heterogeneous across different cortical areas. Indeed, largely due to methodological limitations, the spatiotemporal dynamics of cholinergic activity across the cortex and their relationship with ongoing neural dynamics and behavioral state remain unknown. Here, we combined wide-field "mesoscopic" imaging with expression of a green fluorescent, genetically encoded GPCR-based acetylcholine sensor (GRABACh) and a red fluorescent Ca2+ indicator (RCaMP1b) to simultaneously monitor cholinergic and neural activity across the cortical mantle in awake behaving mice. Neonatal sinus injection of adeno-associated viruses (AAVs) provided widespread expression of both genetically encoded indicators enabling brain-wide mesoscopic imaging in adult animals. We imaged mice during periods of spontaneous behavior while tracking locomotor activity, pupil diameter, facial motion, and neocortical EEG. Our results revealed a robust but spatially distinct coupling between behavioral state and activation of the two sensors. Moreover, the relationship between cholinergic modulation and cortical output varied with state and cortical region. Our results highlight the complex relationship between cholinergic signaling, cortical dynamics, and spontaneous fluctuations in behavioral state and provide novel evidence that state-dependent neuromodulation is not a uniform process, but rather highly variable across the neocortex.

3-011. Weight perturbation outperforming node perturbation

Paul Zuge[1,2] PAUL@ZUEGE-WEB.DE
Christian Klos[1] CKLOS@UNI-BONN.DE
Raoul-Martin Memmesheimer[1,3] RM.MEMMESHEIMER@UNI-BONN.DE

[1]University of Bonn
[2]Institut fur Genetik
[3]Neural Network Dynamics and Computation, Institute for Genetics

Biological constraints often impose restrictions for plausible plasticity rules such as locality and reward-based rather than supervised learning. Two widely studied learning rules that comply with these restrictions are weight (WP) and node (NP) perturbation. They use stochastic perturbations to synaptic weights or neurons to estimate the gradient. NP is considered to be superior to WP as the number of weights and therefore also the perturbation dimension typically massively exceeds the number of nodes in a network.

Here we show that this conclusion no longer holds when we consider two biologically relevant properties: First, tasks extend in time, thus increasing the perturbation dimension of NP but not WP. Second, tasks are low dimensional, with correlated, redundant inputs to neurons. In particular there are many different weight configurations that solve a task. We show that for such cases WP often outperforms NP. We first analytically calculate the expected error curves for linear networks. We find that WP outperforms NP when the effective input dimension is smaller than the number of timesteps in a trial. The analytics further show that when a given target output contains irreproducible parts, this additionally hinders NP but not WP. We also verify that these findings generalize to more realistic settings. Specifically, we show that WP works comparably well or better than NP for an RNN solving a delayed non-match-to-sample task or a layered network solving MNIST. Finally, we find that augmenting WP (or NP) with more evolved optimization techniques further improves its performance.

Taken together, we show that the predominant view that WP performs worse than NP is not true for a large category of biologically relevant cases. This increases the need to investigate WP and treat it as a serious candidate for learning in the brain.

3-012. Probabilistic information processing in humans and recurrent neural networks

Nuttida Rungratsameetaweemana NUTTIDA.RUNGRAT@GMAIL.COM
Robert Kim RKIM@SALK.EDU
Terrence J Sejnowski TERRY@SALK.EDU

Salk Institute for Biological Studies

In nature, sensory inputs are often highly structured, and statistical regularities of these signals can be extracted to form expectation about future sensorimotor associations, thereby facilitating optimal behavior. To date, the circuit mechanisms that underlie these probabilistic computations are not well understood. Through a recurrent neural network (RNN) model and human psychophysics, the present study investigates circuit mechanisms for processing probabilistic structures of sensory signals to guide behavior. We first constructed and trained a biophysically constrained RNN model to perform a probabilistic decision making task similar to paradigms designed for humans. Specifically, the training environment was probabilistic such that one stimulus was more probable than the others. We show that both humans and the RNN model successfully extract information about stimulus probability and integrate this knowledge into their decisions and task strategy in a new environment. Specifically, performance of both humans and the RNN model varied with the degree to which the stimulus probability of the new environment matched the formed expectation. In both cases, this expectation effect was more prominent when the strength of sensory evidence was low, suggesting that like humans, our RNNs placed more emphasis on prior expectation (top-down signals) when the available sensory information (bottom-up signals) was limited, thereby optimizing task performance. Finally, by dissecting the trained RNN model, we demonstrate how competitive inhibition and recurrent excitation form the basis for neural circuitry optimized to perform probabilistic information processing.

3-013. A two-stage model of V2 demonstrates efficient higher-order feature representation

Timothy Oleskiw[1,2] TDO1@NYU.EDU
Ruben Diaz-Pacheco RUIDPR@GMAIL.COM
J. Anthony Movshon[1] MOVSHON@NYU.EDU
Eero Simoncelli[1] EERO.SIMONCELLI@NYU.EDU

[1]New York University
[2]Center for Neural Science

Recently physiology has shown that while single units in area V1 respond primarily to the local spectral content of a stimulus, single units in V2 are selective for higher-order image statistics that distinguish natural images. Despite these observations, a description of how V2 constructs higher-order feature selectivity from V1 outputs remains elusive.

To study this, we consider a two-layer linear-nonlinear network mimicking areas V1 and V2. The V1 stage is built from linear filters that tile the dimensions of position, orientation, and scale, while the V2 stage computes linear combinations of rectified V1 outputs. When connection weights are optimized so that output responses match the higher-order statistics of a texture model (Portilla & Simoncelli, 2000) computed on natural images, the fitted V2-like units resemble localized differences of V1 afferents across all four tuning dimensions. Interestingly, we find these model fits bear strong qualitative resemblance to those fit to data collected from single units in primate V2, suggesting that some V2 neurons are well-suited for encoding these natural image features.

Cortical neurons, such as those of V1, are known to exhibit heavy-tailed (sparse) response distributions to natural images, a fact believed to reflect an efficient image code. Model V2-like units, computing differences over V1 afferents, exhibit a level of sparsity (i.e., kurtosis) similar to what is seen in model V1 populations. In addition, we show that a classifier trained to detect higher-order image features from the kurtosis of responses over space is more efficient when computed on model V2-like units than when computed on comparable V1-like units, in that it requires smaller response ensembles to achieve the same classification accuracy.

Thus, differences over V1 afferent activity are an efficient mechanism for computing higher-order visual features, providing an explanation for the receptive field structures observed in neurons within primate area V2.

3-014. Unsupervised object learning explains face but not animate category structure in human visual cortex

Ehsan Tousi[1,2] E.KTOUSI@GMAIL.COM
Marieke Mur[1] MMUR@UWO.CA

[1]Western University
[2]Psychology

High-level visual cortex represents natural object categories of ecological relevance, including faces and animates. Deep neural networks, which are currently the best computational models of the human visual system, tend to show a weaker natural category structure in their internal representation. An important difference is that humans rely heavily on unsupervised learning during development, while deep neural networks are standardly trained using supervised learning. Recent work using unsupervised learning has shown promise toward providing a better model of human visual processing. Unsupervised models discover structure in their inputs and are not constrained by experimenter-defined categories. We hypothesize that unsupervised learning will give rise to an object representation that emphasizes natural categories and better explains human data. To test our hypothesis, we trained ResNet50 on the ImageNet database using both supervised and unsupervised learning. We characterized the network's internal representation of object images, including faces and animals, over the course of training. To assess the network's ability to predict human data, we correlated its internal representations with the representation of the same images in object-selective visual cortex, measured with functional magnetic resonance imaging in human adults. The unsupervised version of the network better predicts the human object representation than the supervised version. This difference emerges relatively early in training and increases as learning progresses. Better performance after unsupervised training is partly driven by the network's ability to discover natural face category structure in the input images. Importantly, both supervised and unsupervised models fall short of predicting the category division between animate and inanimate objects in the human brain data, suggesting that it is challenging to learn the division from static images alone. Our findings suggest that the natural category structure in human high-level visual cortex may arise from unsupervised learning during development.

3-015. Dynamic cortical network architecture corresponds to state-dependent visual perception

Hadas Benisty[1,2] HADAS.BENISTY@GMAIL.COM
Lavonna Mark[1] LULULAVONNAYILUEN.MARK@YALE.EDU
Jessica Cardin[1] JESS.CARDIN@YALE.EDU
Michael Higley[1] MICHAEL.HIGLEY@YALE.EDU

[1]Yale School of Medicine
[2]Neuroscience

Behavioral states fluctuate dramatically during wakefulness and encompass changes in arousal, locomotion, and general motor activation. Cerebral cortex activity is profoundly modulated by these state transitions, coinciding with variation in the ability to perceive and respond to behaviorally relevant cues. Although the relationship between behavioral state and local circuit dynamics has been well-explored, little is known of how state variables influence large-scale cortical network organization and the consequences for perceptual ability. Moreover, a growing body of work suggests that both single-neuron activity and goal-directed behavior require the coordination of multiple cortical regions, emphasizing the importance of understanding the dynamic functional connectivity of motor, sensory, and association areas during task performance. Here, we use widefield mesoscopic calcium imaging of the entire dorsal neocortex in head-fixed mice performing a visual detection task. In addition to monitoring cortical activity, we quantified spontaneous transitions in behavioral state measured by facial videography, pupillometry, and locomotion on a stationary running wheel. First, we find that increasing arousal levels are associated with significantly enhanced psychometric perceptual ability. Then, using a combination of novel graph analyses for mesoscopic data, we find that the functional connectivity of cortical networks is strongly modulated by state. For spontaneous activity, increasing arousal drives enhanced centrality for visual regions while reducing somatosensory connectivity. Finally, we find that correct task performance on low-arousal trials is associated with network structure preceding trial onset that matches high-arousal epochs. Conversely, network structure on

incorrect trials matches low-arousal epochs. Overall, our findings provide evidence that dynamic fluctuations in large-scale cortical network dynamics provide a unifying mechanism underlying the enhancement of perceptual ability during arousal.

3-016. Attentional modulation of neural activity timescales and their relation to spatial network structure

Roxana Zeraati[1] ROXANA.ZERAATI@UNI-TUEBINGEN.DE
Yanliang Shi[2] SHI@CSHL.EDU
Nicholas Steinmetz[3] NICK.STEINMETZ@GMAIL.COM
Tirin Moore[4] TIRIN@STANFORD.EDU
Tatiana Engel[2] ENGEL@CSHL.EDU
Anna Levina[5] ANNA.LEVINA@UNI-TUEBINGEN.DE

[1]University of Tubingen
[2]Cold Spring Harbor Laboratory
[3]University of Washington
[4]Stanford University
[5]Universtity of Tubingen

Ongoing neural activity unfolds across different timescales reflecting networks' specialization for task-relevant computations. However, it is unknown whether these timescales can be flexibly modulated during trial-to-trial alternations of cognitive states (e.g., attention state) and what mechanisms can cause such modulations. We analyzed autocorrelations of population spiking activity recorded from individual cortical columns of the primate area V4 during a spatial attention task and a fixation task. We estimated timescales from autocorrelations using a novel method based on Approximate Bayesian Computations and applied a Bayesian model comparison to determine the number of timescales in neural activity. We found that at least two distinct timescales are present in both spontaneous and stimulus-driven activity. The slower timescale was significantly longer on trials when monkeys attended to the receptive fields location of the recorded neurons than on control trials when monkeys attended to a different location.

We hypothesized that the observed timescales emerge from the recurrent network dynamics shaped by the spatial connectivity structure. We developed a network model consisting of binary units representing cortical minicolumns with local spatial connectivity among them. We found that the activity of model minicolumns exhibits two distinct timescales: A fast timescale induced by vertical recurrent excitation within a minicolumn and a slow timescale induced by horizontal interactions among minicolumns. The timescales depend on the network topology, and the slow timescale disappears in networks with random connectivity. We derived an analytical relationship between the timescales and connectivity parameters, enabling us to identify model parameters best matching the timescales in the data. The model indicates that modulation of timescales during attention arises from a slight increase in the efficacy of horizontal recurrent interactions. Our results suggest that multiple timescales in local neural dynamics emerge from the spatial network structure and can flexibly adapt to task demands.

3-017. Linking piriform cortex activity to odor perception in the mouse

Pedro Herrero-Vidal[1,2] PMH314@NYU.EDU
Edmund Chong[1] EDMUND.W.CHONG@GMAIL.COM
Cristina Savin[1,3] CS5360@NYU.EDU
Dmitry Rinberg[1] RINBERG@NYU.EDU

[1]New York University
[2]Center for Neural Science
[3]CNS, CDS

The piriform cortex (PCx) is believed to play a key role in creating meaningful perceptual representations of odors, but how it encodes them remains unclear. To investigate this question, we take advantage of recent technological developments that allow us to construct "synthetic odors" by optogenetically activating spatio-temporal patterns of activity in the olfactory bulb (OB). This enables independent and precise control of the neural activity patterns that drive PCx, something unattainable with natural odorants. We record PCx responses in naive mice and measure how they change in response to parametric perturbations of OB activity. We find that PCx neurons are sensitive to spatial and temporal aspects of OB activity, with neural responses affected more strongly when perturbing earlier elements in the OB sequence. We use decoding to assess the behavioral relevance of PCx responses by comparing neural predictions to behavioral outputs in trained animals. Specifically, we train decoders to discriminate between a target and non-target odors, using as input different spatial and/or temporal features of PCx activity.

We compare different PCx codes in terms of their ability to detect the target and their generalization to various perturbations, and compare their predictions to animal behavior in the same task. We find that spatio-temporal codes are best at discriminating the synthetic target odor, matching behavioral discrimination performance. Decoding ability degrades monotonically with the magnitude of the spatial and/or temporal perturbations. Remarkably, the neural predictions closely match the perceptual changes seen in behavior. Overall, our results support the idea that odor identity is encoded in PCx spatio-temporal population activity. Moreover, the fact that PCx responses to perturbations in naive animals can reproduce behavioral reports in trained animals suggests that we are tapping into the intrinsic computations that make PCx a key circuit for the perception of odors.

3-018. Restoring vestibular afferent dynamics improves accuracy of prosthesis-evoked vestibulo-ocular reflex (VOR) responses

Pum Wiboonsaksakul[1,2] PUM@JHU.EDU
Dale Roberts[1] DALE.ROBERTS@JHU.EDU
Charles Della Santina[1] CDS@JHMI.EDU
Kathleen Cullen[1] KATHLEEN.CULLEN@JHU.EDU

[1]Johns Hopkins University
[2]Biomedical Engineering

An exciting and emerging approach to treat patients with impaired vestibular function is a prosthesis that senses head rotation and transforms this movement into vestibular afferent stimulation, substituting for the damaged periphery. Early results from clinical trials, while encouraging, have shown only partial functional improvement. Bridging a gap between basic science knowledge and clinical applications, we implemented biomimetic dynamics in vestibular prostheses for the first time. We asked whether representing the natural dynamics of vestibular afferents in the mapping between head motion and afferent stimulation would result in better performance. To test this proposal, we compared vestibulo-ocular reflex (VOR) responses evoked by the static mapping used by all current devices (no dynamics) to those evoked by 4 newly implemented mappings representing and exceeding the characteristic high-pass dynamics of vestibular afferent processing. Testing was done in two monkeys with profound bilateral vestibular loss that had been implanted with a prosthesis. VOR eye movements were first quantified in response to sinusoidal stimulation that spanned the natural frequency range (0.2 – 20 Hz). We found that afferent-like high-pass mappings evoked more robust VORs with more precise timing. In contrast, the standard static mapping showed a gain decline and increasingly sluggish timing with increasing frequency. Furthermore, mappings with high-pass dynamics exceeding natural range, produced an undesirable phase advance. VOR eye movements were also quantified in response to transient stimulation and similar trends were observed. Overall, using a mapping that mimicked the afferent subclass known to provide primary contribution to the VOR yielded optimal performance. This suggests that endogenous afferent dynamics are well matched to produce accurate VOR response and advocates for a more biomimetic prosthesis design. Together, these results confirm that the implementation of biomimetic mappings in vestibular prostheses can optimize functional outcomes for patients.

3-019. Temporal dynamics of cascaded deep networks

Michael Iuzzolino[1,2] MLIUZZOLINO@GOOGLE.COM
Michael Mozer[3] MCMOZER@GOOGLE.COM
Samy Bengio[4] BENGIO@GOOGLE.COM

[1]University of Colorado Boulder
[2]Computer Science
[3]Google Research, Brain Team; University of Colorado Boulder
[4]Google Research, Brain Team

Deep feedforward neural networks are often considered as a framework for modeling primate vision [1, 2], yet they lack a key property of biological systems, cascaded dynamics. In cascaded dynamics, information propagates from all neurons in parallel but transmission is gradual over time. Cascaded dynamics contrast with the typical operation of feedforward nets in sequential stages wherein each layer fully completes computation before processing begins in subsequent layers. We construct cascaded networks by introducing propagation time delays in deep feedforward nets. We focus on the ResNet architecture, which has skip connections that permit faster transmission of more primitive perceptual representations, leading to an architecture whose functional depth increases over time and yields a trade-off between processing speed and accuracy.

This cascaded architecture is asymptotically equivalent to its feedforward counterpart but exhibits nontrivial temporal dynamics. The equivalence allows us to add temporal dynamics to existing pretrained models. We compare

pretrained-cascaded models with models trained from scratch with a temporal-difference (TD) form of the cross-entropy loss that aims to obtain an accurate response as quickly as possible.

Given truly parallel neural hardware, TD-cascaded nets obtain a strictly superior speed-accuracy profile over feedforward nets run in either a sequential or cascaded manner. TD-cascaded nets are also more robust to various forms of time-varying image modulations, including occlusion and peripheral blurring introduced by fixation shifts, and they more efficiently handle video frame sequences that pan from one object to another. We observe systematic differences in processing efficiency between object classes as well as between instances within a class. TD-cascaded nets classify prototypical examples faster than outliers, with prototypicality assessed by human judgments, clustering methods, and network generalization accuracy. We thus suggest that differences in processing efficiency that have been attributed to cortical feedback processes [e.g., 3, 4] might well be explained by cascaded dynamics.

3-020. Multi-session alignment for longitudinal calcium imaging

Shivin Yadav[1,2] SYADAV@UCSD.EDU
Ryoma Hattori[1] RYHATTORI@UCSD.EDU
Eun Jung Hwang[3] EUNJUNGHWANG.PHD@GMAIL.COM
Asaf Ramot[1] ARAMOT@UCSD.EDU
An Wu[1] ANW103@UCSD.ED
Takaki Komiyama[1] TKOMIYAMA@UCSD.EDU
Gal Mishne[1,4] GMISHNE@UCSD.EDU

[1]University of California, San Diego
[2]Halıcıoğlu Data Science Institute
[3]Rosalind Franklin University
[4]HDSI

Functional imaging (e.g., two-photon calcium imaging) in awake animals has become a standard technique in many systems neuroscience labs. To understand long-term processes such as learning and memory at both neural and behavioral levels, experimental neuroscientists monitor neuronal activity longitudinally while simultaneously recording animal behavior. The ability to record from the same identified population of neurons across days is a major advance, and is expected to raise our understanding of neuronal circuits. One of the crucial post-processing steps of such experiments is the alignment of the recorded imaging data across days to enable one-to-one mapping of neurons across all sessions. This is essential to understanding the changes in neural representation over time. While methods to perform within-session motion correction are common, alignment of two-photon calcium imaging sessions recorded across days is an open problem. Semi-automated approaches exist, but these rely on user input to select matching regions of interest (ROIs) and only align pairs of sessions. Here we present a fast automatic pipeline for multi-session alignment of calcium imaging across days. Our approach is based on classic computer vision algorithms for registration of natural images, which we adapt to multi-session alignment of neuroimaging videos. We validate our approach on several datasets with varying spatial statistics and noise levels, demonstrating fast and accurate imaging alignment and one-to-one mapping of ROIs across multiple sessions.

3-021. Formerly known as latent: Single-trial decision variables recorded from area LIP using neuropixels

Natalie A. Steinemann[1,2] NS3058@COLUMBIA.EDU
Gabriel Stine[1,2] GMS2191@COLUMBIA.EDU
Eric M Trautmann[1] EMT2177@COLUMBIA.EDU
Ariel Zylberberg[3] ARIEL.ZYLBERBERG@GMAIL.COM
Mehdi Sanayei[1] MS4874@COLUMBIA.EDU
Michael N. Shadlen[4] MS4497@COLUMBIA.EDU

[1]Columbia University
[2]Neuroscience
[3]Department of Brain and Cognitive Sciences, University of Rochester, Rochester; Columbia University
[4]Zuckerman Mind Brain Behavior Institute, Columbia Uni.; Kavli Institute; Howard Hughes Medical Inst.

The concept of a dynamic firing rate, r(t), is elemental to principles of computational systems neuroscience. Yet it is elusive, or latent, because single neurons produce spikes punctuated by interspike silent intervals. r(t) is thus estimated by averaging across repetitions on the assumption that the brain has many neurons that also approximate r_bar(t)—what we term the neurophysiologist's assumption of ergodicity (AoE). With few exceptions

these populations have been inaccessible to electrophysiology, either because they lie deep in sulci or because the population is itself heterogeneous. This is especially problematic for the study of stochastic dynamic processes, such as decision making, where a decision variable, V(t) is hypothesized to represent the cumulative sum of iid random values, e_i, as in discrete drift-diffusion. Averaging across repetitions renders only the expectation : the drift rate, or slope of V(t). Novel high-channel count linear probes (primate neuropixel probes; IMEC/HHMI Janelia) now allow us to record large populations of neurons in deep structures, subsets of which appear to represent the same V(t). We can thus observe putative V(t) directly on single trials.

Here we share initial experiences with high density neuropixel recordings from the parietal cortex of rhesus monkeys while they make perceptual decisions. We recorded simultaneously from neurons in the lateral intraparietal area (LIP) and other areas. We observe representations of V_bar(t) in the measured firing rates of as few as 10 neurons on single decisions. The rates exhibit time dependent variance consistent with bounded drift-diffusion. They thus corroborate less direct methods devised to detect these and other properties of LIP neural firing rates and they provide partial support for AoE.

3-022. Predicting perturbation effects from resting state activity using functional causal flow

Amin Nejatbakhsh[1] NEJATBAKHSH.AMIN@GMAIL.COM
Francesco Fumarola[2,3] FRANCESCO.FUMAROLA@RIKEN.JP
Saleh Esteki[4] ME1168@NYU.EDU
Taro Toyoizumi[5] TARO.TOYOIZUMI@RIKEN.JP
Roozbeh Kiani[4] ROOZBEH@NYU.EDU
Luca Mazzucato[6] LMAZZUCA@UOREGON.EDU

[1]Columbia University
[2]RIKEN Center for Brain Science
[3]Center For Brain Science
[4]New York University
[5]RIKEN
[6]Institute of Neuroscience, University of Oregon, USA

Cognitive function arises from large ensembles of interacting neurons. Targeted manipulation of the brain to alter cognitive behavior will be greatly facilitated by understanding the causal interactions within ensembles. A key challenge to achieve such an understanding is that commonly used “large-scale” recording technologies in primates often permit recording from only a small fraction of neurons in a circuit. Are sparse recordings sufficient for inferring causal interactions within a circuit? Are causal interactions inferred from sparsely recorded circuit activity sufficient for predicting the effect of targeted perturbations of neural responses? To address these, we performed a comprehensive theoretical and experimental investigation using causal inference based on delayed embedding. We developed a novel statistical method to infer the causal functional connectivity (“causal flow”) from ensemble recordings of neural activity during the resting state. We first validated our method on synthetic ground truth data from rate networks and biologically plausible models of spiking cortical circuits. We found that causal flow inferred during the resting state reliably predicts the effect of perturbation on network activity, even in the relevant regime of sparse recording. We then applied this method to large-scale recordings from the prefrontal cortex of awake monkeys during quiet wakefulness. We demonstrated that perturbation effects from single-channel micro-stimulation can be accurately predicted from the causal flow inferred during the resting state. These predictions hold at the single electrode level, demonstrating a striking level of granularity and precision. Our results provide a new framework for discovering the rules that enable generalization of resting state causal interactions to more complex behavioral states, paving the way toward targeted circuit manipulations and new brain-machine interfaces.

3-023. Continuous attractor dynamics account for heterogeneity in human behavior during estimation and categorization of integrated sensory evidence

Jose M. Esnaola-Acebes JOSEMESNAOLA@GMAIL.COM
Alex Roxin AROXIN@CRM.CAT
Klaus Wimmer KWIMMER@CRM.CAT

Centre de Recerca Matematica

Perceptual decision and continuous stimulus estimation tasks involve making judgments based on accumulated sensory evidence. Network models of evidence integration usually rely on competition between neural populations each encoding a discrete categorical choice (e.g. whether average stimulus orientation is clockwise (CW) or

counterclockwise (CCW) compared to a reference) [1,2]. By design, these models do not maintain information of the integrated stimulus (e.g. the average stimulus direction in degrees) that is necessary for a continuous perceptual judgement. Here we develop a computational network model that can integrate a continuous stimulus feature such as orientation and can also account for a subsequent categorical choice. The model, a ring attractor network [3], represents the estimate of the integrated stimulus direction in the phase of an activity bump. We reduced the network dynamics of the ring model to a two-dimensional equation for the amplitude and the phase of the bump which allows one to study evidence integration analytically. We found that the model can account for qualitatively distinct temporal integration behaviors, depending on the specific evolution of the amplitude of the bump attractor. These range from early weighting (primacy) over uniform weighting to late weighting (recency), depending on the relative strength of sensory stimuli compared to the amplitude of the bump and also on the initial state of the network. The specific relation between the internal network dynamics and the sensory inputs can be modulated by changing a single parameter of the model, the general excitability of the network. Finally, we show that this model can account for the heterogeneity in how human observers temporally weight evidence across a stream of oriented stimulus frames [4,5]. Overall, our work suggests continuous attractor dynamics as a potential underlying mechanism of stimulus integration and perceptual categorization.

[1] Wang, X.-J. J. Neuron, (2002). [2] Roxin, A., & Ledberg, A. PLoS Comp. Bio., (2008). [3] Ben-Yishai, et al. PNAS, (1995). [4] Wyart, V., et al. Neuron, (2012). [5] Cheadle, S., et al. Neuron, (2014).

3-024. Spiking neural network model of simultaneous localization and mapping with Spatial Semantic Pointers

Nicole Dumont[1,2] NS2DUMON@UWATERLOO.CA
Terry C. Stewart[3] TCSTEWAR@UWATERLOO.CA
Chris Eliasmith[1] CELIASMITH@UWATERLOO.CA

[1]University of Waterloo
[2]Cheriton School of Computer Science
[3]University of Waterloo Collaboration Centre

To navigate in new environments, an animal must be able to keep track of it’s own position while simultaneously creating and updating an internal map of features in the environment, a problem known as simultaneous localization and mapping (SLAM). This requires integrating information from different domains, self-motion cues and sensory information. Recently, Spatial Semantic Pointers (SSPs) have been proposed as a vector representation of continuous space that can be encoded via neural activity. A key feature of this approach is that these spatial representations can be bound with other features, both continuous and discrete, to create compressed structures containing information from multiple domains (e.g. spatial, temporal, visual, conceptual). In this work, SSPs are used as the basis for a biological-plausible SLAM model called SSP-SLAM. It is shown that the self-motion driven dynamics of SSPs can be implemented with a hybrid oscillatory interference/ continuous attractor network of grid cells. The estimated self position represented by this network is used for online learning of an associative memory between landmarks and their positions – i.e. an environment map. This map in turn is used to provide corrections to the path integrator. We show that complex maps can be accurately learned and this use of sensory information results in greatly improved self-position estimation. Furthermore, grid cells, boundary cells, and object vector cells are accounted for in this model.

3-025. Relevance of network topology for the dynamics of neuronal networks with spiking neurons

Simachew Mengiste[1] SIMACHEW@UNI-BONN.DE
Ad Aertsen[2] AERTSEN@BIOLOGIE.UNI-FREIBURG.DE
Arvind Kumar[3] ARVKUMAR@KTH.SE

[1]University of Bonn
[2]University of Freiburg
[3]KTH Royal Inst of Technology

How network structure shapes the activity dynamics of a biological neuronal network is one of the key questions in neuroscience. In rate-based networks, with linear neurons, several features of network structure (e.g. degree correlation, eigenvalue spectra) are correlated with network dynamics and function. However, these results do not directly apply to spiking neuronal networks (SNN). Non-linear properties of neurons, weak synapses and spikes can obscure the relationship between properties of network structure and the dynamics observed in linear rate-based networks. To better understand the structure-dynamics relationship in SNNs, we analysed the descriptors of network structure and activity of over 9,000 SNNs with different sizes and topologies (random, small-world,

scale-free). These networks were degenerated by systematically pruning neurons/synapses according to specific rules. Surprisingly, we found that the topological class was not an indicator of the SNNs activity dynamics as quantified by the firing rate, synchrony and spiking regularity. In fact, we found that networks with qualitatively different topologies could show similar activity dynamics. In most cases, the network activity changes also did not depend on the rules according to which neurons/synapses were pruned from the networks. Our analysis revealed that the effective synaptic weight (ESW) was the most crucial feature in predicting the statistics of spiking activity in SNNs. The ESW is closely related to excitation-inhibition balance and can be experimentally measured for each neuron. We conclude that the network topology and rules by which SNNs degenerate are irrelevant for SNN activity dynamics. Beyond neuroscience, our results suggest that in large networks with non-linear nodes, the ESW among the nodes, instead of the topological network class, may be a better predictor of the network dynamics.

3-026. EqSpike: Spike-driven equilibrium propagation for neuromorphic implementation

Erwann Martin[1] ERWANN.MARTIN@GMAIL.COM
Maxence Ernoult[2] ERNOULT.M@GMAIL.COM
Jeremie Laydevant[3] JEREMIE.LAYDEVANT@CNRS-THALES.FR
Shuai Li[3] SHUAI.LI@CNRS-THALES.FR
Damien Querlioz[4] DAMIEN.QUERLIOZ@UNIVERSITE-PARIS-SACLAY.FR
Teodora Petrisor[5] TEODORA.PETRISOR@THALESGROUP.COM
Julie Grollier[3] JULIE.GROLLIER@CNRS-THALES.FR

[1]Thales TRT
[2]Universite de Montreal, Mila
[3]Unite Mixte CNRS, Thales
[4]Centre de Nanosciences et de Nanotechnologies
[5]Thales

The brain inspiration has driven the design of spiking neural networks demonstrating an outstanding energy efficiency at test time. However, training such networks remains an incredible challenge since applying backpropagation in this context yields non-local computations, which is neither biologically plausible nor hardware friendly. One route is to build on Spike Timing Dependent Plasticity (STDP), a two-factors local learning rule feasible in compact circuitry, which however does not solve the credit assignment problem in deep networks. Another direction is to make backpropagation comply better with locality constraints, often at the cost of adding a third factor gating the weight update and with potential circuit overhead. In this work, we propose a spiking implementation of Equilibrium Propagation (EP), a rate-based algorithm which gets the best of both worlds with a two-factors learning rule being equivalent to backpropagation through time (BPTT) on RNNs fed by static inputs. Our implementation, called EqSpike, results in a learning rule local both in space and time: feedforward activations need not be stored and the gradient computation is fully event-based. We demonstrate EqSpike training on MNIST, achieving 96.9% test accuracy, which is close to rate-based EP and BPTT. We show that EqSpike, if implemented in silicon technology, could cut the energy consumption of inference and training by up to three orders of magnitude compared to Graphical Processing Units (GPUs). Finally, we show that EqSpike weight updates exhibit STDP-like correlations, hilighting its possible connection with biology. Overall, this work demonstrates the benefits of the synergy between neurosciences and neuromorphic hardware design.

3-027. Disparate energy consumption despite similar network activity

Michael Deistler[1,2] MICHAEL.DEISTLER@UNI-TUEBINGEN.DE
Pedro Goncalves[3] PEDRO.GONCALVES@CAESAR.DE
Jakob H. Macke[1] JAKOB.MACKE@UNI-TUEBINGEN.DE

[1]University of Tuebingen
[2]Cluster of Excellence – Machine Learning for Science
[3]research center caesar

Similar neural circuit activity can emerge from disparate sets of parameters such as membrane conductances or synaptic strengths [1]. In turn, the biophysical properties of single neurons can be governed by energy minimisation principles [2]. It is unclear how such principles observed at the single neuron level translate into neural circuits and to what extent metabolic cost constrains the experimentally observed biophysical variability. Here, we combined computational modelling and novel methods for statistical inference to gain insights into the energy consumption of the pyloric network in the stomatogastric ganglion. As previously shown in a model of this circuit,

disparate sets of membrane and synaptic conductances can lead to similar network activity [1]. Further analysing two such circuit configurations with highly similar activity, we found that their energy consumption substantially differs. We then identified all parameter configurations that match experimentally observed activity from the crab Cancer borealis. To do so, we enhanced and used a recently introduced Bayesian inference method based on artificial neural networks [3]. The models we obtained produced activity similar to the experimental data, but their energy consumption differed by almost an order of magnitude. Which membrane and synaptic conductances influence energy consumption? By analysing the space of data-compatible models, we found that some parameters become constrained by requiring energy efficiency, while others barely do. Surprisingly, we found that some conductances can have a large influence on energy despite their currents consuming little energy, and vice versa. Furthermore, we observed that, by varying specific pairs of membrane and synaptic conductances, circuit function can be maintained while energy consumption can be substantially reduced. Our results suggest that, even in complex neural circuits that underlie highly specific functional requirements, some parameter configurations might be metabolically preferable.

3-028. Demixed shared component analysis of neural population data from multiple brain areas

Yu Takagi[1] YUTAKAGI322@GMAIL.COM
Steven Kennerley[2] S.KENNERLEY@UCL.AC.UK
Jun-ichiro Hirayama[3] JUNICHIRO.HIRAYAMA@AIST.GO.JP
Laurence Hunt[1] LAURENCE.HUNT@PSYCH.OX.AC.UK

[1]University of Oxford
[2]University College London
[3]AIST/RIKEN AIP

Recent methodological advances make it possible to record thousands of neurons simultaneously. Although such high-dimensional recording yields insights that are not apparent from studying single neuron activity, analysing population data remains a non-trivial problem because of the heterogeneity of responses and 'mixing' of encoded variables observed in neural data (Fusi et al., 2016). Besides the number of neurons recorded, neuroscientists are also experiencing a revolution in the number of brain areas recorded. Recently, researchers have started to investigate interactions between populations of neurons in distinct areas (e.g., Steinmetz et al., 2019). However, while these approaches yield important insights into cross-regional information sharing, they cannot identify what and when task parameters are shared. It is known that the task parameters are mixed at the level of single neuron (Rigotti et al., 2013) or low-dimensional components (Mante et al., 2013). Thus, it is important to demix which task information is being shared across regions. Here, inspired by a method developed for a single brain area, demixed PCA (dPCA: Kobak et al., 2016), we introduce a new technique for demixing variables across multiple brain areas, called demixed shared component analysis (dSCA). dSCA decomposes population activity into a few components, such that the shared components capture the maximum amount of shared information across brain regions while also depending on task parameters (Fig. 1). Using simulated data (Fig. 2), we show that dSCA finds shared components that are demixed into specific task parameters in a time-resolved manner, whereas previous approaches (CCA/RRR) fail. Next, we reanalyze two previously published datasets (Fig. 3 and 4) and show that dSCA captures shared components among different brain areas specific to decision making task parameters. We believe dSCA will be useful for neuroscientists who will have a large amount of data from different brain areas during complex cognitive experiments.

3-029. Model-based surprise modulates model-free learning: the interaction of novelty, surprise, and reward in human behavior

Alireza Modirshanechi ALIREZA.MODIRSHANECHI@EPFL.CH
He A. Xu HE.XU@EPFL.CH
Marco P. Lehmann MARCO.LEHMANN@GMAIL.COM
Michael H. Herzog MICHAEL.HERZOG@EPFL.CH
Wulfram Gerstner WULFRAM.GERSTNER@EPFL.CH

Ecole polytechnique federale de Lausanne

Humans adapt quickly to changes in the environment, perceived as a feeling of surprise and caused by the violation of expectations. Expectations are formed by world-models learned through exploration, often driven by curiosity. Reinforcement learning (RL) models of human behavior have, however, focused on reward-driven actions, and how surprise influences learning in a reward-based paradigm during exploration and exploitation is still an open question. Here, we propose to integrate surprise and novelty into RL to account for exploration as well as adaptation to environmental changes. We design a deep sequential decision-making paradigm, combine it with modeling, and show that reward, novelty, and surprise are all needed to explain the behavior of human participants. Our experiment contains a long phase of pure exploration where participants take over a hundred actions to find a first reward and a later phase of adaptation caused by an unexpected change in the environment. In our model, adaptive behavior is implemented by a surprise-based control of the learning rate, independently of novelty; and exploratory behavior is implemented by a novelty-seeking strategy, independently of surprise. We find that human decisions are dominated by model-free (MF) action choices but with a learning rate modulated by model-based (MB) surprise. Importantly, even though participants have learned a model of the world, they hardly use it for planning, but mainly to infer if a stimulus is surprising. Our model outperforms its alternatives, accurately predicts the grand majority of individual actions, and allows us to dissociate surprise, novelty, and reward in EEG signals. Whereas notions of novelty, surprise, and reward have each been studied separately before, our computational framework gives, for the first time, a unified formal explanation for how these three drives of behavior are combined to influence action choices and to integrate exploration and adaptation into RL.

3-030. Early phase LTP is sustained through the interplay of spine structural changes and cooperative receptor binding

Moritz Becker[1,2] MORITZ.BECKER@PHYS.UNI-GOETTINGEN.DE
Christian Tetzlaff[3] TETZLAFF@PHYS.UNI-GOETTINGEN.DE

[1]University of Gottingen
[2]III. Institute of Physics – Biophysics
[3]University of Goettingen

Long-term potentiation of synaptic strength (LTP) is one of the main mechanisms underlying memory and learning. The different phases of LTP are subject to various biological and biophysical processes. Understanding these processes, their interactions, and their contribution to the expression of LTP provides important insights into the functioning of learning and the role of synaptic dynamics in neurodegenerative diseases. LTP is regulated mainly by the abundance of AMPA receptors (AMPARs) at synapses. However, it is unclear how the postsynaptic spine regulates AMPAR trafficking to sustain the early phase of LTP (E-LTP) that can last 1-6 hours. We hypothesize that the interplay between specific features of postsynaptic receptor trafficking, notably cooperative receptor binding and exocytosis, is responsible for sustaining synaptic changes during this LTP phase. Therefore, we formalized a biophysical model of AMPAR trafficking that accounts for several experimentally motivated mechanisms and can capture a wide range of experimental findings. In particular, we show that E-LTP is only maintained for hours when receptor trafficking is governed by cooperative receptor binding and spine structural modifications that include spine growth and mobilization of AMPAR-containing recycling endosomes. The correct sequence and timing of these processes are important for E-LTP to occur. This combined model provides a new hypothesis about the principles underlying the maintenance of E-LTP and leads to specific, new experimental predictions that could verify our findings.

3-031. PNS-GAN: conditional generation of peripheral nerve signals in the wavelet domain via adversarial networks

Luke Prince[1] LUKE.Y.PRINCE@GMAIL.COM
Guillaume Lajoie[2] G.LAJOIE@UMONTREAL.CA
Blake Richards[3] BLAKE.RICHARDS@MILA.QUEBEC
Olivier Tessier-Lariviere[1] OLIVIERTL@HOTMAIL.COM
Pascal Fortier-Poisson[4] PASCAL@BIOS.HEALTH
Lorenz Wernisch[4] LORENZ@BIOS.HEALTH
Emil Hewage[4] EMIL@BIOS.HEALTH
Oli Armitage[4] OLI@BIOS.HEALTH

[1]Mila
[2]Universite de Montreal, Mila
[3]McGill Unviersity - Mila
[4]BIOS Health

Simulated datasets of neural recordings are a crucial tool in systems neuroscience for testing the ability of decoding algorithms to recover known ground-truth. Although there has been considerable effort made towards the development of stochastic models of cortical electrophysiology data, models of the typical peripheral nerve recording setups are notably lacking. Since the peripheral nervous system is vital to understanding the intrinsic motivations driving behaviour and the neural basis of chronic disease states, it is important to close this gap. In

this work, we introduce PNS-GAN, a generative adversarial network capable of producing realistic nerve recordings conditioned on physiological biomarkers. PNS-GAN operates in the wavelet domain to preserve both the timing and frequency of neural events with high resolution. PNS-GAN generates sequences of scaleograms from noise using a recurrent neural network and 2D transposed convolution layers. PNS-GAN discriminates over stacks of scaleograms with a network of 3D convolution layers. We find that our generated signal reproduces a number of characteristics of the real signal, including the power spectral density, and more subtle statistics that are consistent with known physiological understanding.

3-032. From single neurons to networks in the dorsal raphe nucleus

Emerson Harkin[1,2] EHARK069@UOTTAWA.CA
Alexandre Payeur[3] APAYE026@UOTTAWA.CA
Jean-Francois Boucher[1] JBOUC096@UOTTAWA.CA
Lea Caya-Bissonnette[1] LCAYA082@UOTTAWA.CA
Dominic Cyr[1] D.CYR18@ROGERS.COM
Andre Longtin[1] ALONGTIN@UOTTAWA.CA
Michael Lynn[1] MLYNN101@UOTTAWA.CA
Chloe Stewart[4] CSTEW017@UOTTAWA.CA
Richard Naud[1] RNAUD@UOTTAWA.CA
Jean-Claude Beique[1] JBEIQUE@UOTTAWA.CA

[1]University of Ottawa
[2]Department of Neuroscience
[3]Universite de Montreal
[4]McGill University

The serotonin neurons of the dorsal raphe nucleus (DRN) are the main sources of serotonergic input to the forebrain and key regulators of learning and behavior. In vivo electrophysiological recordings have shown that these neurons exhibit transient and sustained responses to environmental stimuli, and these distinct responses have been suggested to play different roles in shaping animal behavior. To understand how these response patterns arise, we constructed a detailed bottom-up model of the DRN. Based on a detailed electrophysiological characterization of the two main cell types in the DRN (serotonin neurons and somatostatin-expressing GABAergic interneurons), we developed a novel single-neuron modelling framework that combines the realism of Hodgkin-Huxley models with the simplicity and predictive power of generalized integrate-and-fire (GIF) models. Using a bank of models fitted to individual genetically-identified DRN neurons, we found that simulated spiking networks with a realistic degree of neuron-to-neuron heterogeneity constructed via bootstrapping exhibited a nearly linear input-output transformation compared with homogenous networks. This linear transformation exhibited strongly time-dependent gain with clear transient and sustained components. The gain of the transient component was modulated by a common subthreshold voltage-activated transient potassium current IA which has not previously been implicated in gain control. Finally, we found that the time-dependent gain of DRN output is associated with a pronounced sensitivity to the derivative of external inputs, but that the degree of derivative sensitivity depends on the level of background input. The interplay between background and changing inputs may play a role in shaping the transient and sustained responses of DRN serotonin neurons to environmental stimuli observed in vivo. Together, our results underscore the importance of cellular heterogeneity in determining the functional properties of networks, reveal a new biophysical mechanism of gain modulation, and illustrate their connection to derivative sensitivity in a system of learning and behavioural control.

3-033. Cellular mechanisms for quenching neuronal variability

Chengcheng Huang[1,2] HUANGC@PITT.EDU
Tilo Schwalger[3,4] TILO.SCHWALGER@BCCN-BERLIN.DE
Julia Veit[5,6] JULIA.VEIT@PHYSIOLOGIE.UNI-FREIBURG.DE
Larry Abbott[7] LFABBOTT@COLUMBIA.EDU
Hillel Adesnik[8] HADESNIK@BERKELEY.EDU
Ken Miller[7] KEN@NEUROTHEORY.COLUMBIA.EDU
Brent Doiron[9] BDOIRON@UCHICAGO.EDU

[1]University of Pittsburgh
[2]Neuroscience
[3]TU Berlin
[4]Institute for Mathematics
[5]University of Freiburg
[6]Physiology I

[7]Columbia University
[8]University of California, Berkeley
[9]University of Chicago

A wealth of experimental studies show that the trial-to-trial variability of neuronal activity is quenched during stimulus evoked responses. This fact has helped ground a popular view that the variability of spiking activity can be decomposed into two components. The first is due to irregular spike timing conditioned on the firing rate of a neuron (i.e. a Poisson process), and the second is the trial-to-trial variability of the firing rate itself. Quenching of the variability of the overall response is assumed to be a reflection of a suppression of firing rate variability. Network models have explained this phenomenon through a variety of circuit mechanisms. However, in all cases, from the vantage of a neuron embedded within the network, quenching of its response variability is inherited from its synaptic input. We analyze in vivo whole cell recordings from principal cells in layer (L) 2/3 of mouse visual cortex. While the variability of the membrane potential is quenched upon stimulation, the variability of excitatory and inhibitory currents afferent to the neuron are amplified. This discord complicates the simple inheritance assumption that underpins network models of neuronal variability. We propose and validate an alternative (yet not mutually exclusive) mechanism for the quenching of neuronal variability. We show how an increase in synaptic conductance in the evoked state shunts the transfer of current to the membrane potential, formally decoupling changes in their trial-to-trial variability. The ubiquity of conductance based neuronal transfer combined with the simplicity of our model, provides an appealing framework. In particular, it shows how the dependence of cellular properties upon neuronal state is a critical, yet often ignored, factor. Further, our mechanism does not require a decomposition of variability into spiking and firing rate components, thereby challenging a long held view of neuronal activity.

3-034. Flexible neural coding across behavioral states during a visual change detection task

Alex Piet[1] ALEX.PIET@ALLENINSTITUTE.ORG
Marina Garrett[2] MARINAG@ALLENINSTITUTE.ORG
Anton Arkhipov[1] ANTONA@ALLENINSTITUTE.ORG
Doug Ollerenshaw[2] DOUGO@ALLENINSTITUTE.ORG
Farzaneh Najafi[2] FARZANEH.NAJAFI@ALLENINSTITUTE.ORG
Iryna Yavorska[2] IRYNA.YAVORSKA@ALLENINSTITUTE.ORG
Nick Ponvert[2] NICK.PONVERT@ALLENINSTITUTE.ORG
Peter Groblewski[2] PETERG@ALLENINSTITUTE.ORG
Shawn Olsen[2] SHAWNO@ALLENINSTITUTE.ORG

[1]Allen Institute for Brain Science
[2]Allen Institute

How do representations in sensory cortex respond to changes in visual stimuli, behavioral context, and task relevant goals? To investigate neural coding during behavior we developed a systematic pipeline to train mice to perform a visual change detection task during two-photon calcium imaging. In this task, head fixed mice are shown serially flashed images. When the image identity changes, the mice must lick to receive a water reward. Each mouse is imaged during six consecutive behavioral sessions with changing task demands. Three sessions use a set of familiar images the mice were trained on, followed by three sessions with novel images they have not seen before. Active behavioral sessions on each image set are interleaved with passive behavior sessions where the lick spout is removed and no rewards can be earned. Behavioral quantification shows that individual mice use unique mixtures of several strategies to perform the task including visually driven licking and periodic (timing based) guessing. Further, within each session mice go through epochs of engaged and disengaged behavior. Using calcium imaging of cell-type specific transgenic mice expressing GCaMP6f (pan-excitatory, VIP, and SST) we track individual cells across these diverse behavioral states (n = 90 mice, 539 sessions, 89309 cells). We developed a kernel regression model to fit neural activity and quantify the influence of visual information, behavioral readouts like pupil diameter and running, as well as changing task demands on neural coding across behavioral states. Our model allows quantitative tracking of cell type specific coding strength and dynamics. We find cell type specific changes in feature coding during changes in behavioral state. Our large-scale behavioral and neural dataset and accompanying analysis code will be publicly released in early 2021 and will be of broad interest to the COSYNE community.

3-035. Modeling ephaptic coupling effects in central and peripheral nerve fibers

Helmut Schmidt[1] HSCHMIDT@CBS.MPG.DE
Gerald Hahn[2] GERALD.HAHN@UPF.EDU
Gustavo Deco[2] GUSTAVO.DECO@UPF.EDU
Thomas Knosche[1] KNOESCHE@CBS.MPG.DE

[1]Max Planck Institute for Human Cognitive and Brain Sciences
[2]Universitat Pompeu Fabra

Axonal connections between distant brain areas are regarded as faithful transmitters of action potentials, with axonal microstructure being the sole determinant of their delays. Some studies have challenged this view and found evidence for an effect of ephaptic interaction between axonal connections on spike propagation. A major obstacle is the computational effort involved in such studies, which reduces the focus on few axonal fibers. Here, we propose a novel modeling scheme based on a simplified spike propagation model, which allows us to study ephaptic coupling effects between thousands of axons at reasonable computational cost. We demonstrate that this model can be used to describe ephaptic coupling effects in nerve fibers within the white matter of the brain, as well as in peripheral nerve fibers. The ephaptic interaction between spikes on different axons can be described by coupling functions based on the LFP. Our model reproduces previous findings for the slowing down and synchronization of spike volleys in nerve bundles composed of identical axons. As it is possible to model the interaction between thousands of axons, we are able to study the case of distributed axon diameters, as observed in white matter. We find that the distribution of axon diameters, and the resulting distribution of spike propagation velocities, can tear apart a spike volley and prevent synchronization. In addition, asynchronous spike volleys tend to speed up rather than slowing down due to a net depolarization of the axonal membrane within the spike volley. The magnitude of this effect increases with the number of spikes within a spike volley, which leads to activity-dependent delays. In conclusion, ephaptic coupling effects can now be studied at system-size level, and consequences of activity-dependent delays on flexible and transient synchronization of fast oscillations (such as gamma oscillations) between distant brain areas can be explored.

3-036. Rat sensitivity to multipoint statistics is predicted by efficient coding of natural scenes

Riccardo Caramellino[1,2] RICCARDOCARAMELLINO@GMAIL.COM
Eugenio Piasini[3] EPIASINI@SAS.UPENN.EDU
Andrea Buccellato[4] ANDREA.BUCCELLATO@STUDENTI.UNIPD.IT
Anna Carboncino[1] ACARBONC@SISSA.IT
Vijay Balasubramanian[3] VIJAY@PHYSICS.UPENN.EDU
Davide Zoccolan[1] ZOCCOLAN@SISSA.IT

[1]SISSA Scuola Internazionale Superiore di Studi Avanzati
[2]Cognitive Neuroscience
[3]University of Pennsylvania
[4]Universita di Padova

Extracting useful information from natural images is challenging and requires exploiting their statistical structure. Efficient coding theories posit that this can be done by allocating resources to represent visual features that are more variable, and therefore more informative, across natural scenes. Several studies have shown that, across natural images, there is a precise ranking in the variability of multipoint correlations and that the same ranking is found in human sensitivity to these statistics. This raises the question of whether such specialization of the visual system is hard-wired in the cortical circuitry, as the result of an evolutionary process, or is learned from the statistics of the visual environment during early postnatal development. Testing this hypothesis requires performing invasive experiments in animal models – ideally rodents, which would allow controlled rearing experiments. However, it is unknown whether other mammalian species share with humans the same sensitivity to multipoint correlations. To address this question, we selected four image statistics (from single- to four-point correlations) and trained four groups of rats to discriminate between white noise patterns and binary textures containing variable intensity levels of one of such statistics. An ideal observer model of the task explained well the behavioral performance of individual rats and allowed us to infer their sensitivity to the statistics they were tested on. We found the highest sensitivity for 1- and 2-point correlations, followed by 4-point and finally 3-point correlations. This matches the ranking observed for the variability of these statistics in natural images and for human accuracy at discriminating them from white noise. These results, which were further confirmed by within-subject comparisons of rat sensitivity to pairs of statistics, provide the first demonstration that rats and humans are similarly adapted to process the statistical structure of visual textures

3-037. Integration of identified retinal inputs in wide-field neurons of the superior colliculus

Norma Kuhn[1] NORMA.KUHN@NERF.BE
Chen Li[2] CHEN.LI@NERF.BE
Katja Reinhard[2] KATJA.REINHARD@NERF.BE
Janne Zimmer[2] JANNE.ZIMMER@NERF.BE
Natalia Baimacheva[2] NATALIA.BAIMACHEVA@NERF.BE
Daan Remans[3] DAANREMANS@LIVE.BE
Vincent Bonin[2] VINCENT.BONIN@NERF.BE
Karl Farrow[2] KARL.FARROW@NERF.BE

[1]NERF / KU Leuven
[2]Neuro-Electronics Research Flanders (NERF)
[3]KU Leuven

Tremendous progress has been made in understanding the computational powers of dendrites, from the impact of input location and dendritic morphology to synaptic integration properties. However, how neurons filter specific sensory features along their dendrites to drive behavior is generally unexplored due to the inability to trace the inputs back to their sensory origin from higher brain areas. To explore this question, we use wide-field neurons of the superior colliculus, a genetically targetable cell type that receives direct input from the retina and drives innate defensive behaviors. A combination of transsynaptic circuit tracing, molecular labeling and two-photon calcium imaging was used to identify the retinal inputs to wide-field neurons and measure the response properties both of wide-field neurons and the innervating retinal ganglion cells. Further, by in-vivo measurements of the local signals along the dendrites of wide-field neurons and subsequent linear modeling, we were able to map the retinal inputs onto distinct regions along the dendritic tree. While some dendritic signals can be matched by a linear combination of the retinal inputs, somatic response properties cannot be explained using a linear model. Further, somatic signals were strongly tuned to specific spatial frequencies of a locally structured stimulus of constant mean contrast. This hints to nonlinear processing of the retinal inputs along the dendrites of wide-field neurons and to the importance of including local integration properties into our models of information processing. This unique data set allows us to determine how central neurons combine their specific sensory inputs to selectively respond to behaviorally relevant stimuli.

3-038. Neuromodulators support multiple behaviors through separation of overlapping synaptic memory regimes

Ben Tsuda[1,2] BTSUDA@SALK.EDU
Stefan C Pate[1] SPATE@SALK.EDU
Kay M. Tye[1] TYE@SALK.EDU
Hava T Siegelmann[3] HAVA.SIEGELMANN@GMAIL.COM
Terrence J Sejnowski[1] TERRY@SALK.EDU

[1]Salk Institute for Biological Studies
[2]Computational Neurobiology Laboratory
[3]University of Massachusetts, Amherst

Humans and animals sometimes exhibit drastically different behaviors depending on emotional, arousal, and cognitive state even when presented with identical situations. Such behavioral divergences occur due to differences in underlying neural activity states. Neuromodulatory signaling lies at the core of neural state control, and aberrations result in illness (1). In fact, most psychiatric diseases are characterized by neuromodulatory abnormality, explaining why most psychiatric medications target neuromodulators. Using biologically-inspired artificial neural networks, we show how neuromodulators – modeled as a simple synaptic transmission amplifying/dampening factor – can shift neural networks into new synaptic memory regimes and thereby unlock unique behaviors. We demonstrate how simple neuromodulator-guided shifts in synaptic transmission enable contradictory end behaviors even when presented with identical input stimuli, reminiscent of arousal/mood-modulated behaviors in humans. We then show how a small embedded population of neurons under neuromodulatory control can shift global network dynamics of a neural system, resulting in tailored output behavior. We demonstrate how such systems can be stacked, leading to a myriad of embedded behaviors in response to identical input stimuli. Elaborating this framework of neuromodulation will help our understanding of what goes wrong in diseases, how interventions may ameliorate disorders, and how to build higher capacity, flexible neural networks for machine learning applications.

3-039. Contextual inference underlies the learning of sensorimotor repertoires

James Heald[1,2] JH4069@COLUMBIA.EDU
Mate Lengyel[3,4] M.LENGYEL@ENG.CAM.AC.UK
Daniel M. Wolpert[5] WOLPERT@COLUMBIA.EDU

[1]Columbia University
[2]Neuroscience
[3]University of Cambridge
[4]Department of Engineering
[5]Zuckerman Mind Brain Behavior Institute, Columbia University

Humans spend a lifetime learning, storing and refining a large repertoire of motor memories, as appropriate for the large number of tasks we perform. However, it is unknown what principle underlies how our continuous stream of sensorimotor experience is segmented into separate memories and how we adapt and use this growing repertoire. Here we develop a principled theory of motor learning based on the key insight that memory creation, updating, and expression are all controlled by a single computation – contextual inference – that estimates the probability with which each existing motor memory is appropriate for the current situation. Unlike dominant theories of single-context (Smith et al. 2006; Herzfeld et al. 2014) and multiple-context learning (Haruno et al. 2001; Berniker & Kording 2008; Gershman et al. 2014; Oh & Schweighofer 2019), our repertoire-learning model accounts for key features of motor learning that had no unified explanation: spontaneous recovery (Smith et al. 2006), savings (Kitago et al. 2013), anterograde interference (Sing & Smith 2010), the effect of environmental consistency on learning rates (Herzfeld et al. 2014) and the parcellation of learning into explicit and implicit components (Mcdougle et al. 2016). Critically, our model predicts two novel phenomena – evoked recovery and context-dependent single-trial learning – which we also confirm experimentally. No other theory can unify this broad range of phenomena. These results suggest that contextual inference is the key principle underlying how a diverse set of experiences is reflected in motor behavior, placing the problem of learning a repertoire of memories center stage in motor learning.

3-040. Multiple bumps enhance robustness to noise in continuous attractor networks

Raymond Wang[1] RAYMONDWANG@BERKELEY.EDU
Louis Kang[2] LOUIS.KANG@RIKEN.JP

[1]University of California, Berkeley
[2]RIKEN Center for Brain Science

A central function of continuous attractor networks is encoding continuous variables and accurately updating their values through path integration. To do so, these networks produce localized bumps of activity that move coherently in response to inputs. In the brain, continuous attractors are believed to underlie grid cells and head direction cells that maintain periodic representations of space. However, path integration and various types of periodic tuning can be achieved with any number of activity bumps, and the consequences of producing more or fewer bumps are unclear. To address this problem, we construct continuous attractor networks with different bump numbers and assess their ability to path-integrate in the presence of noise. Single-bump networks with more neurons exhibit greater diffusive motion, which would lead to increased error. However, accuracy can be rescued by introducing multiple activity bumps, which suppresses noise-induced diffusion. Our findings may be particularly salient for large, noisy biological systems, such as the mammalian grid cell network, which could employ multiple attractor bumps to avoid massive error.

3-041. Reconfiguration of sensory cortex and cross-whisker comparisons for discriminating shape

Chris Rodgers[1,2] CCR2137@COLUMBIA.EDU
Ramon Nogueira[1] RN2446@COLUMBIA.EDU
B Christina Pil[1] BUYOUNG.PIL@EINSTEINMED.ORG
Esther Greeman[1] EG3026@COLUMBIA.EDU
Stefano Fusi[1] SF2237@COLUMBIA.EDU
Randy Bruno[1] RB2604@COLUMBIA.EDU

[1]Columbia University
[2]Neuroscience

Humans and other animals can identify objects by active touch, requiring the coordination of exploratory motion and tactile sensation. Both the motor strategies and neural representations employed could depend on the subject's goals. We developed a shape discrimination task that challenged head-fixed mice to discriminate concave from convex shapes. Behavioral decoding revealed that mice did this by comparing contacts across whiskers. In contrast, mice performing a shape detection task simply summed up contacts over whiskers. We recorded populations of neurons in the barrel cortex, which processes whisker input, to identify how it encoded the corresponding sensorimotor variables. Neurons across the cortical layers encoded touch, whisker motion, and task-related signals. Sensory representations were task-specific: during shape discrimination, neurons responded most robustly to behaviorally relevant whiskers, overriding somatotopy. Thus, sensory cortex can be dramatically reconfigured for the task at hand.

3-042. Internal models of task structure shape mesolimbic and dorsomedial striatal dopamine

Marta Blanco-Pozo[1,2] MARTA.BLANCOPOZO@PSY.OX.AC.UK
Thomas Akam[1] THOMAS.AKAM@PSY.OX.AC.UK
Timothy Behrens[1] BEHRENS@FMRIB.OX.AC.UK
Mark Walton[1] MARK.WALTON@PSY.OX.AC.UK

[1]University of Oxford
[2]Experimental Psychology

The world is in continuous change, requiring animals to flexibly adapt their actions to new situations. Dopamine activity is thought to contribute to flexible behaviour by conveying a teaching signal, reward prediction error (RPE), used to update value estimates when predicted and actual outcomes differ. However, a key open question concerns what sources of value information inform dopaminergic RPE's. In particular, it is unclear whether model-based computations that use knowledge of task structure contribute to dopamine activity and whether such signals are broadcast uniformly across projection targets in striatum. To address these questions, we trained mice on a probabilistic multi-step decision making task and recorded calcium activity in dopamine neuron cell bodies and axons, and dopamine release in striatal subregions, using photometry (18 mice, 512 sessions, 184645 trials). Mice chose between two first-step actions which led probabilistically to two second-step states where reward could be obtained. Reward probabilities in the second-step state were anticorrelated. Analysis of both choice behaviour and dopamine activity indicated mice learned this structure; they inferred a single latent variable that controlled both reward probabilities, though, surprisingly, updated this estimate using rewards but not reward omissions. Dopamine responses to the first-step choice in the ventral tegmental area (VTA) and the nucleus accumbens (NAc) comprised both a 'model-free' component that directly reflected previous outcomes following prior choices, and a 'model-based' component that respected both the anti-correlated reward probabilities and transition probabilities mapping first-step actions to second-step states. By contrast, in dorsomedial striatum (DMS) only the model-based component shaped dopamine activity. These signals exist in tandem with influences at different timepoints and timescales from other task factors such as movement, state transition likelihood, and reward rate. Together, these results demonstrate heterogeneous influences on dopamine activity across projection regions.

3-043. Precise inference of single trial network dynamics from calcium imaging using LFADS

Feng Zhu[1] FZHU23@EMORY.EDU
Harrison Grier[2] HARRISONGRIER@UCHICAGO.EDU
Raghav Tandon[3] RAGHAV.TANDON@GATECH.EDU
Changjia Cai[4] CHANGJIA@LIVE.UNC.EDU
Andrea Giovannucci[4,5] AGIOVANN@EMAIL.UNC.EDU
Matthew Kaufman[2] MATTKAUFMAN@UCHICAGO.EDU
Chethan Pandarinath[6] CHETHAN.PANDARINATH@EMORY.EDU

[1]Emory University
[2]The University of Chicago
[3]Georgia Tech
[4]University of North Carolina at Chapel Hill
[5]Joint Department of Biomedical Engineering
[6]Emory University / Georgia Tech

In many brain areas, neural populations act as a coordinated network whose state is tied to behavior on a moment-

by-moment basis and millisecond timescale. Two-photon (2P) calcium imaging is a powerful tool to probe network-scale computation, as it can measure the activity of individual neurons, sample from many layers simultaneously, and identify specific cell types. However, precisely estimating network states from 2P measurements has proven challenging. 2P fluorescence is a noisy, low-pass filtered, nonlinear transformation of neurons' spiking activity, and spiking itself is a noisy reflection of the network's state. Here we describe ZIG-LFADS, a method to precisely infer network states from 2P measurements. ZIG-LFADS extends LFADS - a deep learning method to infer network states from spiking activity - to better model deconvolved calcium signals, whose statistics and temporal dynamics are quite distinct from electrophysiologically-recorded spikes. ZIG-LFADS models deconvolved calcium signals using a zero-inflated gamma (ZIG) distribution, and infers time-varying ZIG parameters for each neuron with the assumption that they reflect an underlying, dynamically-generated network state. We tested our method in two ways. First, we generated synthetic calcium fluorescence traces where the neurons' activity reflected underlying Lorenz systems running at different speeds. Following deconvolution, we compared ZIG-LFADS to both standard LFADS and Gaussian smoothing. ZIG-LFADS inferred latent Lorenz states with higher accuracy than either, particularly for higher-frequency components (>6Hz). We next applied ZIG-LFADS to 2P recordings from 278 modulated caudal forelimb area neurons in head-fixed mice performing a water grab task. ZIG-LFADS substantially improved single-trial decoding of hand position over the other two methods. We then developed a novel strategy to accurately account for the timing of 2P sampling to achieve higher temporal precision, which further improved decoding performance. These results together demonstrate that ZIG-LFADS infers network states from 2P imaging data with unprecedented accuracy and temporal resolution.

3-044. Choice strategies of mice in a stimulus-invariant orientation discrimination task

Dmitry Lyamzin[1,2] DMITRY@BENUCCILAB.NET
Ryo Aoki[3] RYO@BENUCCILAB.NET
Mohammad Abdolrahmani[3] MOHAMMAD@BENUCCILAB.NET
Andrea Benucci[3] ANDREA@BENUCCILAB.NET

[1]RIKEN
[2]Center for Brain Science
[3]RIKEN Center for Brain Science

How the brain computes decision from sensory information is a central question of perceptual decision-making research. To decouple decision variable from the stimulus-specific neural activations, choices can be conditioned on invariant properties of the stimuli. Among visual tasks, orientation discrimination is ideally suited to demonstrate the link between the psychophysical performance and the underlying neural activity. However, it is not clear if mice can learn an invariant orientation discrimination task, and what task execution strategies they would use. Here we present a novel task in which the mouse must report which of the two orientations is more vertical. The same orientation could be a target or a non-target on different trials, making the decision variable invariant with respect to the specific visual stimuli. Using an automated training system, mice (n=40) learned the task reaching a surprising discrimination acuity of up to 6 degrees. To interpret strategies of individual animals, we developed a probabilistic choice model, which allowed for animal-specific biases and variability of orientation estimates, and assumed an optimal combination of noisy sensory information. The model explained variation in performance with task difficulty and predicted a strategy-independent variability due to the circularity of the stimulus space. Furthermore, the model showed that high performance requires high certainty of orientation estimates but certainty for individual stimuli can trade off. Finally, the model revealed a variation in behavior that depended on cognitive state: history-dependent biases had a larger effect on animal's choices during periods of lower engagement in comparison to high engagement. Our results demonstrate that mice are capable of stimulus-invariant discrimination, with a surprising upper bound for discrimination acuity of 6 deg. Animals extracted the decision variable from stimulus-specific visual information often using heuristic strategies that balanced cognitive resources and history-dependent biases.

3-045. Plateau potentials in apical dendrites occur preferentially during unexpected events

Tim M Henley[1,2] HENLEYT@YORKU.CA
Jason E. Pina[1,3] JAYPINA@YORKU.CA
Colleen J Gillon[4] COLLEEN.GILLON@MAIL.UTORONTO.CA
Jerome Lecoq[5] JEROMEL@ALLENINSTITUTE.ORG
Yoshua Bengio[6] YOSHUA.BENGIO@MILA.QUEBEC
Timothy P Lillicrap[7] TIMOTHYLILLICRAP@GOOGLE.COM
Jed Perkins[5] JEDP@ALLENINSTITUTE.ORG
Sam Seid[5] SAMS@ALLENINSTITUTE.ORG
Carol Thompson[5] CAROLTH@ALLENINSTITUTE.ORG
Ryan Valenza[5] RYANV@ALLENINSTITUTE.ORG
Blake A Richards[8] BLAKE.RICHARDS@MCGILL.CA
Joel Zylberberg[1,9] JOEL.ZYLBERBERG@UCDENVER.EDU

[1]York University
[2]Physics
[3]Physics and Astronomy
[4]University of Toronto
[5]Allen Institute for Brain Sciences
[6]Montreal Institute for Learning Algorithms
[7]Google DeepMind
[8]McGill University
[9]Centre for Vision Research

(How) does top-down cortical feedback shape neural activity during learning? Top-down feedback is an essential feature of many computational learning models and predictive hierarchical models of neural representation. Physiological evidence suggests in cortical pyramidal cells top-down feedback may be preferentially targeted to distal apical dendrites and bottom-up input to cell bodies. Recent theoretical work also suggests independent computational roles for apical dendrites and cell bodies in a given neuron, facilitating comparison between top-down predictive signals to bottom-up incoming stimulus signals. Bursting and sustained plateau potentials in the apical dendrites have been suggested to be a crucial factor in plasticity, independent of spike rate. These theoretical studies motivated us to test the hypothesis that top-down feedback, at times of unexpected events, elicits bursting in cortical pyramidal cells. We recorded activity in layer 2/3 (L2/3) and L5 pyramidal neurons separately at their apical dendrites and their somata in the primary visual cortex in mice. Previous work has demonstrated distinct neural responses to expected and unexpected stimuli and suggests these differences are the result of comparison of expected top-down predictions and unexpected bottom-up sensory input that might ultimately drive learning. We recorded the neural response to stimuli with both expected and unexpected statistical properties and compared the changes in these responses in the same dendrites and cell bodies over multiple days. We found the that plateau potentials or bursting occur preferentially at times of unexpected stimulus events in apical dendrites. This difference increased significantly over the multiple days of the experiment. The changes differed between cortical layers as well as between the cell bodies and the dendrites. Our findings support the hypothesis that unexpected events may drive top-down feedback at the apical dendrites.

3-046. Coupled state space models of multi-population recordings

Aneesh Kashalikar[1] ANEESHKASHALIKAR@GMAIL.COM
Joshua Glaser[1] JIG2124@COLUMBIA.EDU
Matthew Whiteway[1,2] M.WHITEWAY@COLUMBIA.EDU
Liam Paninski[1] LIAM@STAT.COLUMBIA.EDU

[1]Columbia University
[2]Zuckerman Institute

It is an exciting time in systems neuroscience, as we can now commonly generate large-scale measurements from multiple brain areas and cell types. There is a great desire to understand the neural dynamics of these emerging datasets - that is, how are populations evolving according to their own internal dynamics, versus being influenced by other populations? State space models, which model high-dimensional neural activity as a low-dimensional representation evolving in time, have been a very successful tool for modeling single populations of neurons. Recent work has begun to apply state space models to multiple populations by using switching linear dynamical system (SLDS) models, in which there are discrete states that dictate the dynamics of all the populations. However, these models do not allow the flexibility for each population to have its own dynamical state, which is likely important for explaining neural data where the populations are not perfectly in-sync.

Here, we present the coupled switching linear dynamical system, which extends SLDS models by providing separate Markov chains of discrete states for each population of neurons. This allows each population to have its own dynamical states. Importantly, the evolution of each populations' dynamical states are not independent, but are coupled within the model. In simulations using this model, we demonstrate that the coupled model can accurately recover the states of each population, unlike a model that does not contain population-dependent states. We then provide a demonstration of this approach on neural data, by fitting the coupled SLDS model to simultaneously recorded neurons in the primary motor cortex and the dorsal premotor cortex. The model, in addition to improving predictive accuracy over baseline models, finds increased inter-region interaction around the beginning of movement onset. Coupled state space models promise to be tools for better understanding the interactions between neural populations.

3-047. Spatial representations rapidly form in parallel with shifting neuromodulatory tone in the hippocampus during novel experience

James Priestley[1,2] JBP2150@COLUMBIA.EDU
Sebi Rolotti[1,2] SVR2111@COLUMBIA.EDU
Jack Bowler[1] JCB2238@COLUMBIA.EDU
Attila Losonczy[1] AL2856@COLUMBIA.EDU

[1]Columbia University
[2]Department of Neuroscience

Principal neurons in the hippocampus exhibit striking selectivity for specific combinations of sensory features, forming representations which are thought to subserve episodic memory. Even during a completely novel experience, ensembles of hippocampal “place cells” are rapidly configured such that the population forms a sparse code over positions; the representation stabilizes within minutes of the first exposure to an environment. How is this rapid encoding of experience achieved? Recent work has implicated a novel “behavioral-timescale” synaptic plasticity (BTSP) rule in hippocampal area CA1 which can rapidly modify neuronal tuning, but it remains unclear how ubiquitous this mechanism is during novel learning.

Here we combine 2-photon calcium imaging with virtual reality in behaving mice, recording simultaneously from hundreds of CA1 neurons during familiar and novel experiences. New place codes developed rapidly when animals were teleported to a novel environment, but the population representation exhibited a transient backward shift in space over this same time period. Contrary to prior reports, this effect was specific to the novel context, and dissipated with familiarity. During this first exposure, many novel place fields exhibited burst firing on their first active laps followed by a backward shift in the location of the place field on later laps, both features characteristic of BTSP.

To probe the mechanisms underlying these transient representational changes, we also recorded calcium dynamics from axonal projections of the locus coeruleus (LC) as they terminate in CA1, using the same context switching paradigm. The first exposure to the novel environment briefly altered the pattern of LC activation, congruent with the hypothesis that broader novelty detection circuits may open a temporary window of heightened plasticity via adjusting neuromodulatory tone. In CA1, we suggest one consequence of this may be to transiently increase the probability of plateau potentials, which could upregulate BTSP during novel experiences.

3-048. A multiscale spiking mushroom body model of learning, memory and motor control in foraging insects

Hannes Rapp[1,2] HANNES.RAPP@ALUMNI.UNI-KOELN.DE
Martin Paul Nawrot[1] MARTIN.NAWROT@UNI-KOELN.DE

[1]University of Cologne
[2]Computatioanl Systems Neuroscience

Foraging is a vital behavioral task for most living organisms. While behavioral strategies have been described in detail for various species, we still lack understanding of the underlying neural circuits and computational principles employed in the nervous system. In this work we shed light on the link between underlying neural circuits and the functional role of computational principles and present how a biologically detailed multiscale neural circuit model of the insect mushroom body (MB) implements sensory processing, learning and motor control. Our work predicts complementary functional roles of population sparse and temporal sparse coding in facilitating memory formation, odor-background segregation, and reliable and fast memory read out under rapidly changing sensory input conditions. We show that realistic dynamic sensory experience leads to distributional shifts in odor identity representations in the MB and that our model can learn a representation that is robust to this non-stationarity.

Finally, we show how knowledge transfer from static to arbitrary complex dynamic conditions can be achieved and allows a cast & surge foraging policy to emerge without explicitly learning it.

3-049. Integrating new memories into the hippocampal network activity space

Giuseppe Pietro Gava[1,2] giuseppepietrogava@gmail.com
Mohamady El-Gaby[3,4] mohamady.el-gaby@ndcn.ox.ac.uk
Stephen McHugh[5] stephen.mchugh@ndcn.ox.ac.uk
Laura Lefevre[5] laura.lefevre@ndcn.ox.ac.uk
Vitor Lopes Dos Santos[5] vitor.lopesdossantos@bndu.ox.ac.uk
Stephanie Trouche[5] stephanie.trouche@igf.cnrs.fr
Simon R. Schultz[1] s.schultz@imperial.ac.uk
David Dupret[5] david.dupret@bndu.ox.ac.uk

[1]Imperial College London
[2]Bioengineering
[3]University of Oxford
[4]Nuffield Department of Clinical Neurosciences
[5]MRC BNDU, University of Oxford

Assimilating new knowledge without corrupting pre-existing memories is a critical brain function. However, learning and memory interact: prior knowledge can proactively influence ongoing learning, and new information can retroactively modify memories of past events [1]. The hippocampus is a brain region essential for learning and memory [2], but the network-level operations that underlie the continuous integration of new experiences into memory, segregating them as discrete traces while enabling their interaction, are unknown. Here we show a network mechanism by which two distinct memories interact. We monitored hippocampal CA1 neuron ensembles in mice as they explored a familiar environment before and after forming a new place-reward memory in a different environment during a conditioned place preference task (CPP). By analysing the activity of principal cells with graph- theoretical techniques, we first found that new associative learning modifies the topology of the cells' co-firing patterns that represents the unrelated familiar environment, suggesting the presence of hysterical dynamics in the network. We further observed that these neuronal co-firing networks evolved along three functional axes: the first identified novelty by segregating the two spatial environments; the second distinguished ongoing behavioural events, and the third revealed cross-memory interaction. We investigated the heterogeneous population of pyramidal cells and found that high activity units quickly distinguished the familiar environment from the CPP apparatus and formed the core representation of each memory. Low activity cells, instead, gradually joined co-activation motifs throughout each task event, enabling cross-memory interactions. Correspondingly, high and low activity principal cells contributed more to the first and third co-firing network axes, respectively, while their interaction mainly supported the second axis. These findings reveal an organizational principle of brain networks where high and low activity cells are differently recruited into coactivity motifs as building blocks for flexible integration and interaction of memories.

3-050. A solution to temporal credit assignment using cell-type-specific modulatory signals

Yuhan Helena Liu[1,2] hyliu24@uw.edu
Stephen Smith[3] stephens@alleninstitute.org
Stefan Mihalas[4] stefanm@alleninstitute.org
Eric Shea-Brown[1] etsb@uw.edu
Uygar Sumbul[3] uygars@alleninstitute.org

[1]University of Washington
[2]Applied Mathematics
[3]Allen Institute
[4]Allen Institute for Brain Science

Animals learn by adjusting the efficacy of their myriad individual synapses, taking into account salient events and neural activity patterns over multiple timescales. How they efficiently solve the underlying temporal credit assignment problem remains elusive. Classical gradient-based methods for training synapses in recurrent neural networks, backpropagation through time (BPTT) and real-time recurrent learning (RTRL), are computationally expensive and violate biological constraints such as locality [1,2]. Meanwhile, biologically plausible approximations to gradient calculations, despite impressive advances, still do not achieve the same learning performance as BPTT and RTRL [1-5]. Here, we show how a previously neglected mechanism can drive better approxima-

tions. Motivated by recent evidence for wide-spread cell-type-specific modulatory signaling [6], we propose a local approximation to gradient calculations, multidigraph-learning (MDGL), that conceptualizes learning via a stack of synaptic and modulatory networks.

Our normative theory posits important roles for neuronal cell types and local diffusive neuromodulation in achieving efficient weight updates. Our learning rule generalizes multi-factor learning, in which a Hebbian eligibility trace is combined with local cell-type-specific modulatory signals in addition to the top-down instructive signals (e.g. dopamine) [7.8]. We trained recurrent spiking neural networks (RSNNs) which obey fundamental biological constraints, including separate excitatory vs inhibitory cell-types and connection sparsity. We observed that our addition improved learning efficiency for several example temporal credit assignment over intervals spanning seconds.

Our learning rule puts forth a novel form of learning shaped by the interplay of synaptic transmission and local cell-type-specific modulatory signals, thereby advancing the field of biologically plausible credit assignment. Moreover, it suggests an efficient on-chip learning algorithm in spike-based hardware for bio-inspired computing. Such approximations of the gradient computation can be especially important as artificial networks become progressively larger and are used to tackle ever more complicated tasks under both time and energy efficiency constraints.

3-051. A shallow artificial neural network recovers structure of visual loom-selective neurons

Baohua Zhou[1] BAOHUA.X.ZHOU@GMAIL.COM
Zifan Li[1] ZIFAN.LI@YALE.EDU
Sunnie S. Y. Kim[2] SUNNIESUHYOUNG@PRINCETON.EDU
John Lafferty[1] JOHN.LAFFERTY@YALE.EDU
Damon Clark[1] DAMON.CLARK@YALE.EDU

[1]Yale University
[2]Princeton University

Animals have evolved sophisticated visual circuits to detect looming objects. Previous studies of the visual system of Drosophila have revealed that looming stimuli are selectively encoded by a suite of neurons, including LPLC2 neurons. Our study aims to understand how the computations in LPLC2 neurons are structured to respond selectively to looming stimuli. We trained an anatomically-constrained shallow neural network to detect whether or not visual signal is an object on a collision course. This enables us to then ask how features of the biological neural circuits relate to the features of the artificial neural network trained on synthetic visual stimuli. In LPLC2 neurons, the dendrites have four branches that extend radially outward in the four cardinal directions and receive motion signals corresponding to those same directions. This distinctive anatomical structure appears to coincide with responsivity to the radial expansion of the edges of a looming object. But is this structured input optimal for detecting looming objects? And how does inference using a population of LPLC2 neurons, which tile visual space, differ from the inference with only a single neuron? To address these questions, we trained a single unit of our model (one LPLC2 neuron) on a set of diverse artificial visual stimuli to detect whether moving objects were on a collision course with the detector location. Surprisingly, there are two opposite solutions: one with dendritic weighting mirroring LPLC2 and the other selective for inward-directed motion. We analyzed how these two solutions each work. Further, we showed that as the population of model units increases, the performance of the model improves, and the biological solution becomes favored. This finding suggests that modeling the entire population of LPLC2 neurons is critical for understanding their encoding of looming stimuli.

3-052. Humans rely on a Mixture Model Learning strategy to solve a multi-task sequential learning problem

Santiago Herce Castanon[1] SANTIAGOHC@GMAIL.COM
Pedro Cardoso-Leite[2] PEDRO.CARDOSOLEITE@UNI.LU
Shawn C. Green[3] CSHAWN.GREEN@WISC.EDU
Paul Schrater[4] SCHRATER@UMN.EDU
Daphne Bavelier[1] DAPHNE.BAVELIER@UNIGE.CH

[1]University of Geneva
[2]University of Luxemburg
[3]University of Wisconsin
[4]University of Minnessota

Humans extract statistical regularities over time by forming internal representations of the transition probabilities between states; in other words, they learn a generative model. Studies on sequence prediction learning typically focus on how experiencing a particular sequence allows the underlying generative model to be learned. Here we ask about the role that generative models themselves play as participants experience and learn to predict different sequences from the same generative family. This novel multi-task prediction paradigm requires participants to go through four sequence learning tasks, with each task being a specific instantiation of a common higher-level generative model. Such a multi-task environment allows participants to display two types of learning: i) within-task learning (that allows them to find the solution to each of the tasks), and ii) across-task learning (which allows them to change the approach that they have for future tasks). To analyze the latter, we focus on the first four trials of each sequence learning task (where behaviour is driven by participants' background knowledge and not by their in-task experience). We show: i) that participants arrive with a structured prior as soon as the first task, and ii) that this prior increasingly resembles the true high-level generative model as participants proceed from one task to the next, showing that participants gain knowledge of the common higher-level generative model. To analyse the former, we turn to behaviour within each task. Here, we model behavior as arising from a mixture model learning strategy, whereby participants have a pool of candidate models (within the space of the true generative model) that they continuously arbitrate against each other to best explain the stimulus sequence. This framework predicts i) the structure in the errors from participants, and ii) a gating mechanism for learning, both of which we then verified in participants' behaviour.

3-053. Geometry of a universal task-specific manifold in the hippocampus

Edward Nieh[1,2] ENIEH36@GMAIL.COM
Manuel Schottdorf[1] MSCHOTTDORF@PRINCETON.EDU
Nicolas Freeman[1] NICOLASWFREEMAN@GMAIL.COM
Sue Ann Koay[1,2] KOAY@PRINCETON.EDU
Lucas Pinto[3,4] LUCAS.PINTO@NORTHWESTERN.EDU
Jeff Gauthier[1] JLGAUTHI@PRINCETON.EDU
Carlos Brody[2] BRODY@PRINCETON.EDU
David Tank[1] DWTANK@PRINCETON.EDU

[1]Princeton University
[2]Princeton Neuroscience Institute
[3]Northwestern University
[4]Physiology

Hippocampal neurons encode physical variables such as space or auditory frequency in cognitive maps. In addition, human fMRI studies have shown that the hippocampus can also encode more abstract, learned variables. However, their integration into existing neural representations of physical variables is unknown. Here, we show the integration of physical and learned task variables as smooth gradients on a single, task-specific, low-dimensional neural manifold, instantiated by the coordinated activity of single neurons. Using 2-photon calcium imaging, we show that individual dorsal CA1 neurons jointly encode accumulated evidence with spatial position in mice performing a decision-making task in virtual reality. Nonlinear dimensionality reduction estimated the neural manifold to be ~4-6 dimensions. Within this low-dimensional space, both physical and abstract variables were jointly mapped in an orderly fashion, creating a geometric representation that we demonstrate to be similar across animals. The existence of conjoined cognitive maps suggests that the hippocampus performs a general computation – to create geometric representations of learned knowledge instantiated by task-specific low-dimensional manifolds.

3-054. Physiological correlates of perceptual threshold

Mitchell Morton[1,2] MITCHELL.MORTON@YALE.EDU
Sachira Denagamage[1] SACHIRA.DENAGAMAGE@YALE.EDU
Isabel Blume[1] IBLUME@COLGATE.EDU
John Reynolds[3] REYNOLDS@SALK.EDU
Monika Jadi[1] MONIKA.JADI@YALE.EDU
Anirvan Nandy[1] ANIRVAN.NANDY@YALE.EDU

[1]Yale School of Medicine
[2]Neuroscience
[3]Salk Institute for Biological Studies

Little is known about the physiological underpinnings that cause correct ('hit') or incorrect ('miss') responses at perceptual threshold. We hypothesized that fluctuations in physiological states in the sensory neocortex determine

whether a trial at threshold will result in a hit or a miss. We analyzed cortical layer-specific electrophysiological recordings in visual area V4 collected while monkeys performed a cued attention task that required them to detect an orientation change. We find that hit trials are characterized by a larger pupil diameter and lower incidence of microsaccades, indicative of increased arousal and greater perceptual stability. Threshold stimuli evoke elevated multiunit activity in hit trials compared to miss trials, across all cortical layers. Broad spiking neurons in the superficial cortical layers exhibit lower variability in hit trials. Hits are also characterized by greater interlaminar coherence between the superficial and deep layers in the pre-stimulus period, and between the input layer and both the superficial and deep layers in the stimulus-evoked period. Taken together, these results indicate that elevated levels of arousal and perceptual stability along with privileged processing of sensory stimuli contribute to hits at perceptual threshold.

3-055. A real-time spike sorting software for thousand electrodes

Pierre Yger PIERRE.YGER@INSERM.FR
Baptiste Lefebvre BAPTISTE.LEFEBVRE@INSERM.FR
Giulia LB Spampinato GIULIA.SPAMPINATO@INSERM.FR
Olivier Marre OLIVIER.MARRE@INSERM.FR

Institut de la Vision, Sorbonne Universite

Recent technological advances have made it possible to record simultaneously from tens to thousands electrodes packed with high density. To analyze these extracellular data, scalable, accurate and semi-automated algorithms have been proposed to sort spikes from hundreds of recorded cells (Pachitariu et al. 2016, Yger et al. 2018). However, these algorithms do not allow tracking the activity of individual neurons during the experiment, since the entire processing is run offline. This is a limitation for experiments where some decisions of the experimentalist could be guided by the recent activity of the recorded cells, and more generally for the design of closed loop experiments. To address this issue we designed an online spike sorting software that accurately sorts spikes in real time for up to a thousand of electrodes. Our algorithm identifies the template waveforms and their spike times by combining a clustering algorithm and a greedy template matching algorithm. It handles well-known spike sorting issues such as misalignments in the spike detection or overlapping spike waveforms. The online clustering procedure allows dealing with slow changes of the templates over time, due to slow drifts of the electrodes. Depending on the number of electrodes to process, it can be run on a few desktop computers that communicate together through ethernet ports. We validated that our software could sort spikes in real time for an increasing number of electrodes on simulated datasets. We assessed the accuracy of the results with both in vitro and in vivo ground truth datasets. Our software thus enables optimal experimental design and closed loop experiments, where the choice of the stimuli to present can be made as a function of the data acquired recently. It will also be a valuable tool for experimentalists to monitor their large-scale recordings during the experiment.

3-056. Decision-making as a lens for identifying animal learning rules

Zoe Ashwood[1] ZASHWOOD@PRINCETON.EDU
Nicholas Roy[1] NROY@PRINCETON.EDU
Ji Hyun Bak[2] JIHYUN.BAK@UCSF.EDU
The International Brain Laboratory[3] INFO@INTERNATIONALBRAINLAB.ORG
Jonathan Pillow[4] PILLOW@PRINCETON.EDU

[1]Princeton University
[2]University of California, Berkeley
[3]The International Brain Laboratory
[4]Princeton Neuroscience Institute, Princeton University

How do animals learn to perform new tasks? What are the rules governing learning? These are open questions in neuroscience. Whereas cognitive science often takes a top-down approach to modeling human and animal behavior [Griffiths, 2010], here we introduce methods for directly inferring learning rules from animal training data. Our method efficiently infers the trial-to-trial changes in an animal's policy, and decomposes those changes into a learning component and a noise component. This allows us to: (i) compare different learning rules and objective functions that an animal may be using to update its policy; (ii) estimate distinct learning rates for different parameters of an animal's policy; (iii) identify variations in learning across cohorts of animals; and (iv) uncover trial-to-trial changes that are not captured by normative learning rules. After validating our framework on simulated data, we applied our model to data from rats and mice learning binary perceptual decision-making tasks. We considered variants of the policy gradient learning rule known as REINFORCE [Williams, 1992], and found that certain rules were far more capable of explaining the trial-to-trial policy changes used by real animals. Whereas the average contribution of the conventional REINFORCE rule to the policy updates used by mice was just 30%, adding reward-offset (baseline) parameters allowed the learning rule to explain 92% of the animals' updates. Intriguingly, the best-fitting learning rates and baseline values indicate that an animal's policy update, at each trial, does not occur in the direction that maximizes expected reward. Understanding how an animal transitions from chance-level to high-accuracy performance when learning a new task not only provides neuroscientists with insight into their animals (and may allow for faster, more efficient training), but also provides concrete examples of biological learning algorithms to the machine learning community.

3-057. Is the brain macroscopically linear? A system identification of resting state dynamics

Erfan Nozari[1] ERFAN.NOZARI@UCR.EDU
Jennifer Stiso[2] JENI.STISO@GMAIL.COM
Lorenzo Caciagli[2] LCACIAGL@SEAS.UPENN.EDU
Eli Cornblath[2] ELI.CORNBLATH@PENNMEDICINE.UPENN.EDU
Xiaosong He[2] HEXS@SEAS.UPENN.EDU
Maxwell Bertolero[2] MB3152@SEAS.UPENN.EDU
Arun Mahadevan[2] ASM7@SEAS.UPENN.EDU
George Pappas[2] PAPPASG@SEAS.UPENN.EDU
Danielle Bassett[2] DSB@SEAS.UPENN.EDU

[1]University of California, Riverside
[2]University of Pennsylvania

A central challenge in the computational modeling of neural dynamics is the trade-off between accuracy and simplicity. At the level of individual neurons, nonlinear dynamics are both experimentally established and essential for neuronal functioning. One may therefore expect the collective dynamics of massive networks of such neurons to only increase in their complexity, thereby supporting the assumption that an "accurate" model of whole-brain dynamics must be nonlinear. To what extent this assumption holds, however, has remained an open question. Here, we provide a rigorous and data-driven answer by investigating whole-brain blood-oxygen-level-dependent (BOLD) and macroscopic field potential dynamics and leveraging the theory of system identification. Using functional magnetic resonance imaging (fMRI) and intracranial electroencephalography (iEEG), we model the spontaneous, resting state activity of 700 subjects in the Human Connectome Project (HCP) and 122 subjects from the Restoring Active Memory (RAM) project using state-of-the-art linear and nonlinear model families. We assess relative model fit using predictive power, computational complexity, and the extent of residual dynamics unexplained by the model. Contrary to our expectations, linear auto-regressive models achieve the best measures across all three metrics, eliminating the trade-off between accuracy and simplicity. To understand and explain this linearity, we highlight four properties of macroscopic neurodynamics which can counteract or mask microscopic nonlinear dynamics: averaging over space, averaging over time, observation noise, and limited data samples. Whereas the latter two are technological limitations and can improve in the future, the former two are inherent to aggregated macroscopic brain activity. Our results demonstrate the discounted potential of linear models in accurately capturing macroscopic brain dynamics. This potential, together with the unparalleled interpretability of linear models, can greatly facilitate our understanding of macroscopic neural dynamics, which in turn may facilitate the principled design of model-based interventions for the treatment of neuropsychiatric disorders.

3-058. A statistical framework for extracting time-varying oscillations from neural data

Andrew Song[1] ANDREW90@MIT.EDU
Francisco Flores[2,3] FJFLORES@NEUROSTAT.MIT.EDU
Demba Ba[4,5] DEMBA@SEAS.HARVARD.EDU
Emery Brown[1] ENB@NEUROSTAT.MIT.EDU

[1]Massachusetts Institute of Technology
[2]Massachusetts General Hospital and Harvard Medical School
[3]Anesthesia, Critical Care and Pain Medcine
[4]Harvard University
[5]School of Engineering and Applied Sciences

To study the neural basis of brain states, such as states of consciousness during the induction of anesthesia, and how they relate to interactions between brain regions, it is important to track the dynamics of spectral power and time-domain quantities, such as phases, in specific frequency bands of the neural data. Bandpass filtering offers a simple and popular approach to this problem. However, the absence of a corresponding statistical generative

model makes it difficult to quantify the uncertainty of power or phase estimates and, as a consequence, to assess the statistical significance of the subsequent inferences. Approaches to mitigate these limitations partition time series into independent, approximately stationary intervals, and fit a stationary generative model to each interval. The assumption of independence of successive intervals results in distortion/discontinuity around the boundaries, which can lead to erroneous scientific interpretation and phenomena such as the “phase slip”. We propose a statistical generative model, termed the piecewise locally stationary oscillatory (PLSO) model, that decomposes a time series with slowly-varying spectra into piecewise-stationary oscillatory components. PLSO models the signal dynamics for the entire data, preventing artifacts from data partitioning, and providing estimates and uncertainty for time-domain quantities of interest. In addition, PLSO can capture spectral dynamics in the time-frequency domain. We first apply PLSO to rat hippocampal LFP data from an open field task. By leveraging PLSO’s ability to conduct inference on the entire time series, we show that the coupling of population spikes to theta phase observed in previous studies is indeed scientifically significant. Moreover, PLSO successfully identifies the prominent spectral components, specifically the theta band, and denoises the signal without introducing boundary artifacts. We also apply PLSO to EEG from humans under propofol anesthesia and show that PLSO can capture smooth spectral dynamics while removing noise/motion artifacts.

3-059. Information-limiting correlations linearize the neural code

Sacha Sokoloski[1,2] SACHA.SOKOLOSKI@MAILBOX.ORG
Ruben Coen-Cagli[3] RUBEN.COEN-CAGLI@EINSTEINMED.ORG

[1]Albert Einstein College of Medicine
[2]Systems and Computational Biology
[3]Albert Einstein College Of Medicine

In the framework of probabilistic population codes, information processing in neural circuits is defined by Bayesian inference, and a population code is "linear" when neurons implement optimal Bayesian inference with linear operations. Empirical work has shown that local, Bayesian computations in neural circuits can often be modelled with simple linear operations, and theoretical work has shown that a neural code is linear as long as the neural activity satisfy certain constraints. Nevertheless, it remains unclear if neural activity typically satisfy these constraints, and therefore if neural codes in the brain are truly linear.

For neurons that fire independently and with Poisson-distributed spike counts, neural codes are linear when the average, total activity of the neural population does not vary with the stimulus — in other words, when the sum of their tuning curves is constant. It has been suggested that the sum of neural tuning curves should approach a constant as the number of neurons in a population increases, such that linear codes naturally arise in large populations of independent neurons. In simulations of orientation-tuned neurons, however, we find that this is not the case, and rather that there is a large performance gap between optimal codes and approximate, linear codes.

To address this performance gap, we first generalize the constraints on neural activity to include a simple bias-correction term. We then add sensory noise to our population models, which induces information-limiting correlations between the neurons. We find that in sufficiently large populations of neurons with information-limiting correlations, bias-corrected linear codes achieve near-optimal performance. Our results suggest a functional role for information-limiting correlations, namely that redundancy in a neural code can increase its linearity.

3-060. Distinct synaptic plasticity mechanisms determine the diversity of cortical responses during behavior

Michele Insanally[1,2] MNI1@NYU.EDU
Badr Albanna[1] BADR.ALBANNA@PITT.EDU
Brian DePasquale[3] DEPASQUALE@PRINCETON.EDU
Saba Fadaei[4] SABA.SHOKATFADAEI@NYULANGONE.ORG
Kanaka Rajan[5,6] KANAKA.RAJAN@MSSM.EDU
Robert Froemke[7] ROBERT.FROEMKE@NYULANGONE.ORG

[1]University of Pittsburgh
[2]Department of Otolaryngology
[3]Princeton University
[4]New York University, School of Medicine
[5]Mount Sinai, Icahn School of Medicine
[6]Department of Neuroscience
[7]New York University

Spike trains recorded from the cortex of behaving animals can be complex, highly variable from trial to trial, and therefore challenging to interpret. A fraction of cells exhibit trial-averaged responses with obvious task-related features such as pure tone frequency tuning in auditory cortex. However, a substantial number of cells (including cells in primary sensory cortex) do not appear to fire in a task-related manner 1,2 and are often neglected from analysis. A recent study presented at Cosyne 2020 used a novel single-trial, spike-timing-based analysis to show that both classically responsive and non-classically responsive cortical neurons contain significant information about sensory stimuli and behavioral decisions suggesting that non-classically responsive cells may play an underappreciated role in perception and behavior 3 (Fig 1). Here we expand this investigation to explore the synaptic origins and potential contribution of these cells to network function. To do so, we trained a novel spiking recurrent neural network (RNN) model that incorporates spike-timing-dependent plasticity (STDP) mechanisms to perform the same task as behaving animals (Fig 2). By leveraging excitatory and inhibitory plasticity rules this model reproduces neurons with response profiles that are consistent with previously published experimental data, including classically responsive and non-classically responsive neurons. We found that both classically responsive and non-classically responsive neurons encode behavioral variables in their spike times as seen in vivo. Interestingly, heterosynaptic plasticity in excitatory-to-excitatory synapses increased the proportion of non-classically responsive neurons and may play a significant role in determining response profiles. Finally, our model also makes predictions about the synaptic origins of classically and non-classically responsive neurons which we compare to in vivo whole-cell recordings taken from the auditory cortex of behaving animals. This approach successfully recapitulates heterogeneous response profiles measured from behaving animals and provides a powerful lens for exploring large-scale neuronal dynamics and the plasticity rules that shape them.

3-061. Fragmented spatial maps from surprisal and affordances

Mirko Klukas[1,2] MIRKO.KLUKAS@GMAIL.COM
Sugandha Sharma[1,3] SUSHARMA@MIT.EDU
Ila R. Fiete[1] FIETE@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Department of Brain and Cognitive Sciences
[3]Brain and Cognitive Sciences

When animals explore spatial environments, their representations often fragment into multiple maps. What determines these map fragmentations, and can we predict where they will occur with simple principles?

We pose the problem of fragmentation of an environment as one of online spatial clustering. We develop a theory in which fragmentation decisions are driven by surprise, by navigational affordances, and by path integration error. When these criteria are implemented in various environments, they reproduce map fragmentations seen in navigating animals. Augmented with a spatial memory, the theory accounts for the reuse of map fragments in environments with repeating substructures. A grid- and place-cell network model with boundary vector cells can implement these abstract criteria to generate appropriate fragmentations.

Existing models of map fragmentation fall into two categories: The first assumes that fragmentation is mainly driven by large path integration errors combined with similarity of sensory observations. This implies that in an environment with no sensory ambiguity there would be no grid remapping. The second considers eigenvectors of different types of transition matrices, e.g. eigenvectors of the successor representation or the graph Laplacian of the adjacency matrix. These models require a full buildup of the transition matrix, and it is not clear how remapping decisions would be driven in a novel environment upon first visit because state transition information has yet to be collected.

In contrast, our model is capable of both, it remaps in environments without sensory ambiguity and also in novel environments upon first visit. Moreover, unlike eigenvectors of transition matrices but consistent with experiments, the grid cells in our model maintain the same internal structure and relationships even as their spatial tuning changes across environments.

Finally, we show that fragmentation permits accurate representation of large spaces in which path integration errors would otherwise dominate.

3-062. Decoding lingual-palatal contacts from population responses in primate orofacial sensorimotor cortex

Derrick Tang[1,2] DTANG7@UCHICAGO.EDU
Antony Simonoff[3,4] ASIMONOFF@UCHICAGO.EDU
Jd Laurence-Chasen[1] JDLAURENCE@UCHICAGO.EDU
Fritzie Arce-McShane[1] FRITZIEA@UCHICAGO.EDU

[1]University of Chicago
[2]Organismal Biology and Anatomy
[3]The University of Chicago
[4]Department of Neurobiology

Complex orofacial sensorimotor behaviors, such as feeding and speech, rely on precise control from the orofacial sensorimotor cortex (oSM). The ways in which tactile and proprioceptive feedback contribute to this process have yet to be characterized. Here, we evaluated the presence of information on lingual-palatal contacts in oSM and how this was affected by absent tactile inputs. Neuronal activity from primary motor (M1) and somatosensory (S1) regions of oSM was recorded simultaneously with 3D tracking of tongue kinematics while monkeys engaged in natural feeding. Nerve blocks to sensory branches of CNV knocked out tactile inputs from the palate, teeth, and tongue while preserving proprioceptive inputs. We implemented two K-Nearest Neighbor (KNN) classifiers to evaluate the ability of spiking activity of M1 and S1 neurons to predict 1) contact of any part of the tongue with six palatal regions encoded by a sequence of six bits (total of 64 possible combinations) and 2) contact of tongue-tip with two anterior palatal regions. For decoding multiple lingual-palatal contacts, classifier performance was significantly better than chance for all cortical regions, with M1 being the most reliable predictor of contact with and without tactile feedback (Fig. 1). The removal of tactile inputs degraded performance when using spiking activity from areas 1/2. In contrast, performance improved with areas 3a/3b. For single lingual-palatal contact, the classifier correctly identified palatal regions above chance with and without nerve block (Fig. 2). The weak effect of absent tactile inputs on decoding lingual-palatal contact suggests alternative sources of sensory information, possibly from proprioceptive inputs from the tongue. The robust representations of lingual-palatal contacts in oSM suggest that their cortical control plays an important role in vital and critically important functions such as feeding. The results may have important implications for the evaluation and treatment of sensorimotor disorders affecting the orofacial system.

3-063. Event-driven spiking neural networks for pattern recognition

Antoine Grimaldi[1,2] ANTOINE.GRIMALDI@UNIV-AMU.FR
Victor Boutin[3] VICTOR.BOUTIN@UNIV-AMU.FR
Laurent U. Perrinet[3] LAURENT.PERRINET@UNIV-AMU.FR
Sio-Hoi Ieng[4] SIO-HOI.IENG@UPMC.FR
Ryad Benosman[4] RYAD.BENOSMAN@UPMC.FR

[1]Insitut de Neuroscience de la Timone
[2]NeOpTo team
[3]Inst. Neur. Timone, Aix-Marseille Univ, Marseille, France
[4]Institut de la Vision - Sorbonne Universite

We propose a neuromimetic Spiking Neural Network (SNN) architecture able to perform pattern recognition. To achieve this, we extended an existing event-based algorithm Lagorce et al. 2017 which introduced novel spatio-temporal features: time-surfaces. Built from asynchronous events acquired by a neuromorphic camera, these time-surfaces allow to code the local dynamics of a visual scene and to create an efficient hierarchical event-based pattern recognition architecture. Our work is three-fold. First, through the analysis of this existing method we were able to adapt its formalism in the computational neuroscience domain by demonstrating it may be implemented using a SNN of leaky integrate-and-fire (LIF) models and Hebbian learning. Then, the classification was performed using multinomial logistic regression and we validated it with a more complex and widely used dataset (N-MNIST, Orchard et al. 2015). A significant contribution was to achieve online classification of the N-MNIST dataset reaching state of the art performances. The third contribution came by adding a homeostatic regulation rule to Hebbian learning as inspired by biological findings. According to the efficient coding hypothesis, neural activity should be equally distributed between neurons. We used this principle to force neurons within the same layer to spike on average with balanced firing rates by setting for each neuron a gain depending on their past activity. Efficiency of this technique was demonstrated through an improvement in spatio-temporal patterns which were learned during the training phase and a classification performance reaching 97, 8% accuracy. As a summary, we were able to develop a neuromimetic SNN model for online digit classification. We aim at pursuing the study of this architecture for natural scenes and hope to offer insights on the efficiency of neural computations.

3-064. Comparative analysis of optogenetic modulation of neural activity in mouse and monkey visual cortex

Alessandro Sanzeni[1] ALESSANDRO.SANZENI@DUKE.EDU
Agostina Palmigiano[2,3] AP3676@COLUMBIA.EDU
Junxiang Luo[4] JULUO@SALK.EDU
Jonathan Nassi[4] NASSI@SNL.SALK.EDU
John Reynolds[4] REYNOLDS@SALK.EDU
Mark Histed[5] MARK.HISTED@NIH.GOV
Ken Miller[2] KDM2103@COLUMBIA.EDU
Nicolas Brunel[1] NB170@PHY.DUKE.EDU

[1]Duke University
[2]Columbia University
[3]Center for Theoretical Neuroscience
[4]Salk Institute for Biological Studies
[5]National Institute of Mental Health

Primary visual cortex (V1) neurons in the mouse and the monkey have some shared response properties, such as tuning for edge orientation, but some differences, such as dramatically different firing rate changes to optimal stimuli. To understand potential differences in circuit properties in the two species, we examined optogenetic responses in mice [Histed 2018] and monkeys [Nassi et al. 2015] to combinations of visual stimuli and optogenetic excitation of pyramidal cells.

In both species, the distribution of optogenetic modulations of visual responses was broad, with means that were much smaller than the standard deviation and with large fractions of suppressed cells. The main quantitative differences between the two species were that responses to visual stimuli, their optogenetic modulations and the fraction of suppressed cells were larger in monkeys. The heterogeneity of optogenetic modulations of visual responses underlies two simple patterns at the population level: (1) In monkeys, at all contrasts, optogenetic stimuli strongly modulated the activity of single neurons, without significantly modifying the distribution of rates across the population. This feature was specific to optogenetic stimuli (visual stimuli shifted the distribution toward higher rates) and appeared also in mice at high contrast. (2) Cells with similar visual responses in the two species had similar statistics of optogenetic modulation. Consistent with this last observation, a network model with the same underlying connectivity reproduces the statistics of optogenetic modulations observed in the two species, provided visually-driven inputs in mice are weaker than in monkeys.

Our work provides the first analysis of how optogenetic stimuli modulate visual responses in mice and monkeys. It suggests that, despite the different way in which the visual cortex represents information, optogenetic modulations of visual responses in the two species follow a similar underlying principle.

3-065. Balanced networks and belief propagation: Stable, responsive and precise

Vincent Bouttier[1,2] VINCENT.BOUTTIER@HOTMAIL.FR
Renaud Jardri[3] RENAUD.JARDRI@CHRU-LILLE.FR
Sophie Deneve[1] SOPHIE.DENEVE@GMAIL.COM

[1]Ecole Normale Superieure
[2]Cognitive Science Department
[3]University of Lille

Humans are reasonably good at dealing with probabilities. Psychophysical experiments demonstrate that humans are close to ideal Bayesian observers. One therefore may wonder how Bayesian inference is carried out in the brain. A possible implementation is that the brain contains networks whose structure mirrors probabilistic graphs and that the network activity corresponds to Belief Propagation algorithm (BP) in the underlying graph [Shon2005, Steimer2009]. BP computes marginals of the probability distribution by sending probabilistic messages throughout the graph. A striking property of BP is that despite the recurrent nature of probabilistic graphs, messages going in opposite directions are decorrelated thanks to the removal of redundant information, thus avoiding "positive feedback" and wrong computations. The corresponding neural network would need to be micro-balanced, with each recurrent excitatory input controlled by an inhibitory input of similar strength [Deneve2016, Ahmadian2019]. Such tight balance is unlikely to be perfectly achieved in brain networks. An imperfect control would introduce deviations from exact inference in human observers, possibly forming aberrant beliefs and modified behavior. This lack of precise control presumably relates to excitation-inhibition imbalance in the brain, a biological marker of autism and schizophrenia [Sohal2019]. We thus study the network dynamics in the case where the removal process is defective [Jardri2013]. First, we observe that impaired inhibitory control destabilizes the network by shifting its eigenvalues. Secondly, the network becomes overconfident: the generated beliefs are too strong given the evidence entering the network. Lastly, the impairment provides undesired memory to the system, which can destroy its ability to adjust rapidly to new situations. Altogether, our work shows that removal of redundant information (due to reciprocal connections) by balancing inhibition allows the system to be precise and flexible, and stable, expanding the set of probability distributions which can be used, in particular those with strong correlations between variables.

3-066. Distributed sampling-based Bayesian inference in coupled neural circuits

Wenhao Zhang[1] WENHAOZ@UCHICAGO.EDU
Tai Sing Lee[2] TAI@CNBC.CMU.EDU
Si Wu[3] SIWU@PKU.EDU.CN
Brent Doiron[1] BDOIRON@UCHICAGO.EDU

[1]University of Chicago
[2]Carnegie Mellon University
[3]Peking University

The brain performs probabilistic inference to interpret the external world, but the underlying neuronal mechanisms remain poorly understood. The stimulus structure of natural scenes exists in a high-dimensional (HD) feature space, and how the brain represents and infers the joint posterior distribution in this rich, combinatorial space is a challenging problem. There is added difficulty when considering the neuronal mechanics of this representation, since many of these features are computed in parallel by distributed neural circuits. Here, we present a novel solution to this problem. We study continuous attractor neural networks (CANNs), each representing and inferring a stimulus attribute, while attractor coupling supports sampling-based inference on the multivariate posterior of the high-dimensional stimulus features. Using perturbative analysis, we show that the dynamics of coupled CANNs realizes Langevin sampling on the stimulus feature subspace embedded in neural population responses. In our framework, feedforward inputs convey the likelihood, reciprocal connections encode the joint stimulus priors, and the internal Poisson variability of neurons generates the correct random walks for sampling. Our model achieves high-dimensional joint probability representation and Bayesian inference in a distributed manner, where each attractor network infers the marginal posterior of the corresponding stimulus feature. The stimulus feature can be read out simply with a linear decoder based only on local activities of each network. In sum, our study provides insight into the fundamental neural mechanisms for realizing efficient HD probabilistic inference.

3-067. The story of a recurrent neural network's training from initialization to task mastery

Owen Marschall[1,2] OEM214@NYU.EDU
Kyunghyun Cho[1] KYUNGHYUN.CHO@NYU.EDU
Cristina Savin[1,3] CS5360@NYU.EDU

[1]New York University
[2]Center for Neural Science
[3]CNS, CDS

Recent years have seen an influx of competing proposals for biologically plausible learning in recurrent neural networks, but it is challenging to compare these models with neural data. One reason is that, experimentally, we have access to only neural activity at scale, not synaptic strengths. Therefore we can assess how dynamics change as an organism gains mastery of a task, but not how the synaptic strengths themselves change. The two are deeply related, but we lack theoretical tools to study synaptic changes through the lens of neural dynamics. In this work we use RNNs to examine this relationship. With new topological methods, one can open the 'black box' of a fully trained RNN and reverse-engineer its function from a dynamical systems perspective (Sussillo & Barak, 2013; Maheswaranathan et al., 2019). Here we generalize these tools to track the progress of an RNN over its learning trajectory. What might this look like? Does the network beeline to the end solution, methodically building up the necessary features (fixed points, line attractors, etc.)? Or are there growing pains, moments when it stumbles in and out of partial solutions before finding an effective mechanism? What about when it fails–is it because the network never develops useful features in the first place, or because it can't arrange them properly? Do the answers change for different learning algorithms? We address these questions in the context of a simple memory task for two example algorithms. For successful training, our method reveals a latent set of learning stages, not apparent from the network's performance alone nor naive geometric notions of change, while unsuccessful training does not get the network even past the first stage. This topology-based analysis of

learning shows promise as a tool for exploring the complex relationship between neural plasticity and dynamics.

3-068. Emergence of tuned inhibitory synapses as a result of interaction between inhibitory plasticity rules

Fereshteh Lagzi[1,2] LAGZI@UCHICAGO.EDU
Martha Canto Bustos[3] MBC27@PITT.EDU
Anne-Marie Oswald[3,4] AMOSWALD@PITT.EDU
Brent Doiron[1] BDOIRON@UCHICAGO.EDU

[1]University of Chicago
[2]Department of Neurobiology
[3]University of Pittsburgh
[4]Neuroscience

It is known that inhibition is an active source of decorrelation between excitatory neurons. Also, in the balanced state, excitation and inhibition are tightly correlated such that the activity of the network as a whole is in the asynchronous state with neurons emitting spikes with irregular timing. The role of different inhibitory subtypes in this picture is largely unexplored. Because the connectivity pattern of inhibition can substantially change network dynamics, it is important to study the role of distinct inhibitory subtypes in shaping correlated spiking activity. To understand the emergence of inhibitory circuits, we use in vitro patch clamp techniques and cell-specific optogenetic stimulation to study the evolution of synaptic weights from different inhibitory subtypes in response to repeated activation. More specifically, we measure the spike timing dependent plasticity (STDP) rules of the inhibitory inputs onto excitatory (E) pyramidal neurons from both Parvalbumin (PV) and Somatostatin (SOM) interneurons in mouse orbital frontal cortex (OFC). Motivated by these experimental findings, we use network modeling to show that the interaction between different inhibitory plasticity rules during learning can result in the emergence of tuned PV to E connections. In our model, SOM neurons develop lateral inhibition which correlates excitatory neurons within cell assemblies, however, PV neurons decorrelate excitatory neurons in a selective manner- to provide network stability. We predict that in associative cortices, if SOM neurons are strongly connected to pyramidal cells and to PV cells, then the PV cells will develop tuned connections to the excitatory cells. Our prediction is supported by theoretical analysis of a plastic network composed of different inhibitory subtypes with plasticity rules consistent with experimental findings. In sum, the plasticity rules of distinct inhibitory subtypes are such that E and PV neurons coordinate themselves into shared assemblies, that are modulated by a separated group of SOM neurons.

3-069. Primitive object perception during larval zebrafish escape behavior

Andrew Bolton[1] ANDREW.D.BOLTON@GMAIL.COM
Olivia McGinnis[2] OLIVIAJMCG@GMAIL.COM
Vikash Mansinghka[3] VKM@MIT.EDU
Florian Engert[1] FLORIAN@MCB.HARVARD.EDU

[1]Harvard University
[2]University of Oxford
[3]Massachusetts Institute of Technology

Perception and characterization of solid objects in our physical world is critical to our interaction with our surroundings and a central feature of intelligence. To probe the evolutionary origins of physical knowledge, we sought to study whether a primitive organism, the larval zebrafish, is sensitive to the locations and properties of obstacles in its environment. The zebrafish, which has become a useful model to study brain-wide circuit dynamics during behavior, engages in reflexive fast escape swims to avoid predators. At the neural level, the escape behavior is executed by a well-studied circuit called the brainstem escape network (BEN), which induces large-angle escape turns to either the right or left. We posited that collisions with solid objects during escape from predators would be maladaptive to the zebrafish. Therefore, obstacle avoidance should play a role in the fish's escape strategy, and the direction of BEN-induced turns should be informed by the locations of barriers. In order to test this hypothesis, we designed a closed-loop high-speed imaging rig outfitted with barriers. We show that zebrafish integrate three-dimensional form, color, and distance to robustly bias their escapes away from solid objects in range of their escape swims. Moreover, through two-photon laser ablation and calcium imaging of neurons in the BEN, we identify a novel inhibitory input to the BEN from the visual system that implements the escape bias we observe. In light of our results, we suggest that zebrafish implicitly construct “object solidity” via simple visuomotor rules that are constructed through modulation of evolutionarily ancient predator-escape circuits.

3-070. Whence motor preparation?

Marine Schimel[1,2] MMCS3@CAM.AC.UK
Ta-Chu Kao[1] TCK29@CAM.AC.UK
Guillaume Hennequin[1,3] G.HENNEQUIN@ENG.CAM.AC.UK

[1]University of Cambridge
[2]Engineering Department
[3]Department of Engineering

In ballistic movement tasks, the primary motor cortex (M1) systematically exhibits preparatory activity prior to movement onset, even in tasks with no enforced delay periods. These preparatory events are largely understood as seeding the dynamics of the movement epoch with movement-specific initial conditions. While computational models show how such initialization might be achieved, they also make clear that movement could instead rely on input-driven M1 dynamics—supported by recent experiments in mice—with no need for preparation at all. Whence motor preparation, then?

Here, we revisit the role of preparatory processes by asking when and why they might emerge as optimal control strategies. We consider the joint dynamics of a recurrent neural network (RNN) model of M1 coupled to a model arm, and use a standard control cost functional to quantify performance in a delayed-reaching task. We use the iLQR algorithm to find the spatiotemporal patterns of network inputs that minimize this cost functional. Critically, these inputs can arise both before and during movement; thus, our framework accomodates a continuum of motor strategies going from purely autonomous motor generation following preparation, to purely input-driven unprepared dynamics.

We find that preparation-based strategies emerge robustly in several RNN architectures, with optimal inputs arising well before movement begins. Interestingly, the benefits of preparation are most pronounced for an inhibition-stabilized network model of M1, which was shown previously to capture quantitative aspects of monkey M1 activity during reaching. In contrast, more weakly connected networks do not benefit as much from preparation. Finally, we propose a new computational method to reveal the features of a recurrent network architecture that make motor preparation beneficial. This meta-learning approach is based on differentiating through optimal control solutions, and could lead to novel hypotheses about the nature of M1 dynamics.

3-071. Contributions of molecular ion channel noise to macroscopic spike time variability

Shusen Pu[1,2] SXP600@CASE.EDU
Peter Thomas[3] PJTHOMAS@CASE.EDU

[1]Vanderbilt University - Case Western Reserve University
[2]Biomedical Engineering
[3]Case Western Reserve University

The contributions of independent noise sources to the variability of action potential timing has not previously been studied at the level of individual directed molecular transitions within a conductance-based model ion-state graph. The underlying connection provides an important example of how mathematics can be applied to study the effects of unobservable microscopic fluctuations to macroscopically observable quantities. We study a stochastic Langevin model and show how to resolve the individual contributions that each transition in the ion channel graph makes to the variance of the interspike interval (ISI). We extend the mean–return-time (MRT) phase reduction developed in (Cao et al. 2020, SIAM J. Appl. Math) to the second moment of the return time from an MRT isochron to itself. Because fixed-voltage spike-detection triggers do not correspond to MRT isochrons, the interphase interval (IPI) variance only approximates the ISI variance. We find the IPI variance and ISI variance agree to within a few percent when both can be computed. Moreover, we prove rigorously, and show numerically, that our expression for the IPI variance is accurate in the small noise (large system size) regime; our theory is exact in the limit of small noise. By selectively including the noise associated with only those few transitions responsible for most of the ISI variance, our analysis extends the stochastic shielding (SS) paradigm (Schmandt et al. 2012, Phys. Rev. Lett.) from the stationary voltage-clamp case to the current-clamp case. We show numerically that the SS approximation has a high degree of accuracy even for larger, physiologically relevant noise levels. We show that the ISI variance is not an unambiguously defined quantity, but depends on the choice of voltage level set as the spike-detection threshold, both in vitro and in silico.

3-072. A thalamo-orbitofrontal circuit signals the control of learning rate

Vijay Mohan K Namboodiri[1,2] VJLBYM@GMAIL.COM
Taylor Hobbs[3] THOBBS13@UW.EDU
Ivan Trujillo Pisanty[4] PISANTY.IVAN@GMAIL.COM
Rhiana Simon[3] RHSIMON@UW.EDU
Madelyn Gray[3] MMGRAY@UW.EDU
Garret Stuber[3] GSTUBER@UW.EDU

[1]University of California, San Francisco
[2]Neurology
[3]University of Washington
[4]Langara College

Learning to predict rewards is essential for the sustained fitness of animals. Contemporary views suggest that such learning is driven by a reward prediction error (RPE) — the difference between received and predicted rewards. The magnitude of learning is proportional to the product of the RPE and a learning rate. Here we demonstrate using two-photon calcium imaging and optogenetics in mice that certain functionally distinct subpopulations of orbitofrontal cortex (OFC) neurons signal the suppression of the learning rate of a reward, either due to reward prediction in a stationary context, or due to the presence of a highly salient aversive stimulus in a context containing unpredicted rewards. We demonstrate that our results are parsimoniously explained by learning rate control. Consistent with a control of learning rate, trial-by-trial fluctuations in OFC activity positively correlates with behavioral learning when RPE is positive, and negatively correlates when RPE is negative. Overall, we present fourteen major findings and rule out ten alternative models. Lastly, using combined imaging and optogenetics, we demonstrate that medial thalamic inputs to OFC sculpt OFC reward response adaptation reflecting learning rate control. These results support emerging theoretical views that the prefrontal cortex encodes and controls learning parameters (i.e. meta learning).

3-073. What is the spatial scale of sleep oscillations in cortex?

Maryam Hasanzadeh Mofrad[1,2] MHASANZA@UWO.CA
Ana Suller Marti[1] ANA.SULLERMARTI@LHSC.ON.CA
Lyle Muller[3] LMULLER2@UWO.CA

[1]Western University
[2]Department of Applied Mathematics
[3]University of Western Ontario

In the two-stage model of memory consolidation, memories are first formed in the hippocampus and then transferred to neocortex for long-term storage. This transfer is mediated by synaptic plasticity, for which the co-occurrence of pre- and post-synaptic activity plays a critical role. While it has become increasingly clear that sleep oscillations actively and causally contribute to this process, it remains unclear to what extent these oscillations coordinate activity across areas in neocortex. Based on early recordings in animals under anesthesia, sleep-like rhythms were generally considered to occur across a wide range of cortex, creating a state of large-scale synchrony. Recent work, however, has challenged this idea by reporting isolated sleep rhythms such as local slow-oscillations and spindles. What is the spatial scale of sleep rhythms in cortex? To answer this question, we adapted deep learning algorithms initially developed for detecting earthquakes and gravitational waves in high-noise settings to analysis of neural recordings in sleep. We studied sleep spindles in non-human primate electrocorticogram (ECoG), human electroencephalogram (EEG), and clinical intracranial electroencephalogram (iEEG) recordings. Our approach, which detects a range of clearly formed large- and small-amplitude spindles during sleep, reveals that sleep spindles co-occur over a broad range of cortex. In particular, multi-area spindles, where more than 10 electrode sites exhibit this rhythm simultaneously, are much more frequent than previously estimated by amplitude-thresholding approaches, which tend to select only the highest-amplitude spindles and can miss events that temporarily fall below threshold. Because spindles are intrinsically related to sleep-dependent consolidation of long-term memories, these results have important implications for the distribution of engrams in cortex of primate and humans. We lastly present results from human EEG under varying memory loads and show our method detects an increase in multi-area spindles following high-load visual working memory tasks.

3-074. Bidirectionally connected cores in a mouse connectome

Jun Kitazono C-KITAZONO@G.ECC.U-TOKYO.AC.JP
Yuma Aoki AOKI-YUMA711@G.ECC.U-TOKYO.AC.JP
Masafumi Oizumi C-OIZUMI@G.ECC.U-TOKYO.AC.JP

The University of Tokyo

It has been suggested that feed-forward processing alone is insufficient to generate conscious experiences, and recurrent (bidirectional) processing is necessary. Therefore, to understand the mechanism of consciousness, it is important to evaluate which parts of the brain network are strongly interconnected in a bidirectional manner.

Here we propose a method for extracting subnetworks that consist of strong bidirectional connections. First, we define a "core" of a network as a subgraph that consists of stronger bidirectional connections than any of its supergraphs. The strength of the connection is evaluated by a graph cut that considers bidirectionality of connections. We then show that, because of two properties of graph-cut weight, submodularity and monotonicity, a network is uniquely decomposed into a nested hierarchical structure by these cores. A larger core includes a smaller core with stronger connections. The largest core, which is the entire network, has the weakest connections and the locally smallest cores have the locally strongest connections. This hierarchical structure allows us to find cores in polynomial time, as opposed to the exponential time it would take to find them by brute force according to the definition of cores.

We applied the proposed method to a whole-brain mouse connectome and show that the cores with strong bidirectional connections consist of brain areas that have been suggested to be highly relevant to consciousness (e.g., areas in isocortex and thalamus, and claustrum). In contrast, these cores do not contain areas thought to be irrelevant to consciousness (e.g., areas in cerebellar cortex and cerebellar nuclei). To evaluate the importance of bidirectionality to the results, we also performed the same analysis for the network where the direction of connections is omitted. We show that when the direction is omitted, cores do not necessarily consist of the brain areas related to consciousness.

3-075. Learning-related dynamics across multiple timescales in cortico-recipient midbrain neurons

Tai-Ying Lee TAI-YING.LEE@DPAG.OX.AC.UK
Andrew King ANDREW.KING@DPAG.OX.AC.UK
Johannes Dahmen JOHANNES.DAHMEN@DPAG.OX.AC.UK

University of Oxford

Perception is shaped not only by sensory information but also by experiences, indicating sensory processing involves integration of both external inputs and internal information. While much research studies the flow of sensory information through feedforward sensory pathways, comparatively little is known about how top-down internal states influence early processing stages of the sensory systems. Even less is known about how such influence changes over the course of learning.

Here, we report preliminary findings from tracking the same population of neurons over the course of learning in one of the major feedback circuits in the mouse auditory system. Mice were trained to perform an auditory detection task, while two-photon calcium imaging was used to selectively measure the activity of neurons in the dorsal cortex of the inferior colliculus (DCIC) that receive input from the auditory cortex (AC). Nonnegative tensor decomposition revealed a set of low-dimensional components that describe the common within-trial temporal dynamics of each neuronal subpopulation, as well as how these patterns were differentially activated across trials and days of learning. Learning-related trial-to-trial variability was found to associate with not only stimuli but also behavioural variables such as whether a reward was acquired. Overall, our analysis suggests that subpopulations of cortico-recipient DCIC neurons show rich learning-related dynamics associated with various task variables across multiple timescales.

3-076. The learning dynamics of automatic cross-domain abstraction and generalisation of latent rules in non-human primates

Michał Wojcik[1] MICHAL.WOJCIK@PSY.OX.AC.UK
Dante Wasmuht[1] DANTE-WASMUHT@HOTMAIL.COM
Makoto Kusunoki[2] MAKOTO.KUSUNOKI@PSY.OX.AC.UK
Miki Kadohisa[1] MIKI.KADOHISA@PSY.OX.AC.UK
Nicholas Myers[1] NICHOLAS.MYERS@PSY.OX.AC.UK
Laurence Hunt[1] LAURENCE.HUNT@PSYCH.OX.AC.UK
John Duncan[3,1] JOHN.DUNCAN@MRC-CBU.CAM.AC.UK
Mark Stokes[1] MARK.STOKES@PSY.OX.AC.UK

[1]University of Oxford
[2]Department of Experimental Psychology, University of Oxford
[3]University of Cambridge

Flexible behaviour requires the brain to extract abstract representations of the rules that govern the environment and generalise them to novel contexts. To study the signatures of these abstract representations, we recorded spiking data from two macaques performing a passive object- association task (Fig1). In phase 1, the animal was presented with a cue and a target, a non-linear combination of which predicted reward (Fig1B). In phase 2, a second set of stimuli was introduced to test whether the rule learned in the first phase cross-generalised to the new sensory domain (Fig1C). Additionally, in both phases, the target had two features: one relevant (shape) and one irrelevant (width) for reward prediction. Theoretical work indicates that solving such a task requires the PFC to generate a specific geometrical architecture: (1) the PFC collapses across features that carry similar behavioural meaning (H1: abstraction, Fig2A), and (2) these newly constructed abstract dimensions generalise across different sensory instances (H2: generalisation, Fig2B). Our study investigated how these processes arise as a function of learning, even in the absence of goal- directed behaviour. We found neuronal markers of learning: the PFC activity predicted reward and collapsed across dimensions that were not predictive (as indicated by a representational similarity analysis (RSA) template-based regression; Fig3). Crucially, in phase 2, an abstract representation of the cue's meaning (context) developed across training sessions (H1 tested by the context model, Fig4A). Consistent with H2, a discriminant function trained to differentiate the cues in the first set of stimuli cross-generalised to the second, novel set (Fig4B). Our results indicate that learning was resolved across training sessions (Fig5A) and that it was accompanied by the emergence of an abstract contextual signal (Fig5B).

3-077. Task-related suppression of correlated variability in auditory cortex predicts behavioral performance but not changes in neural discrimination

Charles Heller[1,2] HELLERC@OHSU.EDU
Daniela Saderi[1] DANIELA.SADERI@GMAIL.COM
Stephen David[1] DAVIDS@OHSU.EDU

[1]Oregon Health and Science University
[2]Otolaryngology

Sound-evoked activity of single neurons in primary auditory cortex (A1) is modulated by active task engagement, enhancing the neural discriminability of task-relevant sound features. At the neural population level, task engagement also suppresses correlated variability (i.e., noise correlations) between pairs of simultaneously recorded neurons. Theoretical work suggests that a reduction in noise correlations can improve sensory coding accuracy, but it is unknown how these task related changes actually impact coding in A1 or behavioral performance. We analyzed data from A1 of ferrets performing a go / no-go tone-in-noise detection task. Multi-channel neural data was collected using linear microelectrode arrays during alternating passive and active (engaged) listening blocks. As expected, transitions from passive to active listening were associated with modulation of evoked activity and decreased noise correlation. Using an optimal linear decoder, we found that discrimination of task-relevant sound categories by neural populations was improved during active listening blocks. Notably, this effect could be almost entirely attributed to modulation of single neuron responses; discrimination did not benefit from reduced noise correlations. Furthermore, we found that noise correlation strength could predict trial-trial behavior performance, while neural decoder accuracy could not. These results suggest that downstream brain areas utilize a non-optimal readout of activity in A1.

3-078. Spontaneous traveling waves are an intrinsic feature of cortical dynamics that impact perception

Zachary Davis[1,2] ZDAVIS@SALK.EDU
Gabriel Benigno[3,4] GBENIGNO@UWO.CA
Terrence J Sejnowski[1] TERRY@SALK.EDU
John Reynolds[1] REYNOLDS@SALK.EDU
Lyle Muller[5] LMULLER2@UWO.CA

[1]Salk Institute for Biological Studies
[2]Systems Neurobiology Laboratory
[3]Western University
[4]Applied Mathematics
[5]University of Western Ontario

Neocortical activity is highly variable, fluctuating from moment-to-moment, with broadly distributed spectral energy (Shadlen and Newsome 1998). In multielectrode array recordings made in area MT of the common marmoset, we find that these fluctuations are neither fully random nor synchronous across space: rather, they are often spatiotemporally organized into traveling waves. As they traverse the cortex, these waves modulate spontaneous and stimulus-evoked spiking activity and perceptual sensitivity as measured in monkeys trained to detect faint visual targets. To gain insight into the neural mechanisms underlying traveling waves, we study a large-scale spiking network model with conductance-based synapses, biologically realistic topographic connectivity and action potential propagation speeds consistent with those observed in unmyelinated horizontal fibers (Girard et al. 2001). We find that these properties are sufficient to generate spontaneous waves across the entire range of network parameters that produce asynchronous-irregular spiking dynamics (Brunel 2000, Renart et al. 2010). Further, we find that neuronal participation in these waves is sparse, enabling traveling waves to coexist with asynchronous-irregular spiking activity without necessarily inducing correlations, which have been found to impair perception (Nandy et al. 2019). This sparse-wave network regime remained sensitive to feed-forward driving input and modulated the strength of stimulus-evoked responses in a manner similar to that observed in cortex. This was in contrast to smaller scale networks that produce dense spike coupling to traveling waves, which drove strong pairwise correlations and rendered the network insensitive to feed-forward driving input. We conclude that the presence of traveling waves in cortical dynamics can be explained as a natural consequence of activity propagating along unmyelinated horizontal fibers. Traveling waves, which improve perceptual sensitivity in MT (Davis et al. 2020), appear to be an intrinsic feature of cortical dynamics, and they therefore likely impact the moment-to-moment processing of information throughout the brain.

3-079. Neuronal morphology influences spike generation and excitability type

Robert Gowers[1,2] GOWERSRO@HU-BERLIN.DE
Susanne Schreiber[1] S.SCHREIBER@HU-BERLIN.DE

[1]Humboldt-Universitat zu Berlin
[2]Institute for Theoretical Biology

Neurons can exhibit different spiking onset types which influence how the input is encoded and - due to their close association with a neuron's phase response curve (PRC) - indicate how weakly connected neurons may synchronize. Conductance-based neuron models seek to capture the different spiking onset types via the bifurcation associated with the onset of regular firing. Previous research, however, has largely considered neurons as point-like and abstracted the neuron's morphology. Using a spatial model that includes morphological parameters, we show that if a neuron's conductance is increased via its morphology, spiking onsets differ from those of a point-neuron model with an equivalent conductance.

To address the effect of morphology on the spiking onset bifurcation, we analysed an active conductance-based soma attached to a passive dendrite. Previous studies have found that increasing the leak conductance can switch the spiking onset from type I to type II. In our spatial model, we increased the input conductance by widening the dendritic diameter. In comparison with a point neuron, the switch from type I to type II occurs at a higher input conductance. Meanwhile, onset bifurcations indicating bistability of spiking and resting states, such as homoclinic and limit-cycle fold bifurcations, occur at higher conductances with narrower conductance ranges. These changes manifest in both the bifurcation structure and PRCs, and are attributable to the lower passive input impedance of the spatial model at higher frequencies.

In addition to the analysis above, we simulated f-I curves and PRCs of anatomically-detailed morphological reconstructions from NeuroMorpho.Org, which illustrate good agreement with our reduced dendrite-and-soma approximation. Overall, we show that the spike generation type depends on dendritic morphology, priming neurons with extensive dendritic arborizations for specific neuronal dynamics, network computations, and possibly contributing to differential variability of different brain areas to neuropathologies such as epileptic seizures.

3-080. A mesoscale connectome defines low-dimensional dynamics of neural activity across the mouse cortex

Yanliang Shi SHI@CSHL.EDU
Tatiana Engel ENGEL@CSHL.EDU

Cold Spring Harbor Laboratory

Massively-parallel neurotechnologies enabled recordings of neural activity on the brain-wide scale, as well as provided detailed maps of the brain-wide anatomical connectivity. The large-scale datasets reveal rich spatiotemporal patterns of neural activity, which are highly variable from trial-to-trial and widely distributed, with behaviorally relevant signals spread across many brain regions. The activity correlations between regions (functional connectivity) closely correspond with the region-to-region anatomical connectivity, suggesting that widespread neural dynamics could arise from simple mechanistic interactions between recurrently connected brain areas. However, the principles defining these brain-wide interactions remain unknown. We developed a computational framework for integrating the connectome with mesoscopic functional neural dynamics across the mouse cortex. We combined the datasets of inter-regional anatomical connectivity and the widefield calcium imaging of neural activity across the mouse dorsal cortex during decision making and spontaneous activity. We found that on single trials, the activity of each cortical area was accurately predicted by the input from other areas, defined as a sum of their activities weighted by the anatomical connection strengths. Analyzing the composition of inputs received by each area, we identified six cortical areas (driver areas) that provided the main input to all other areas (driven areas). Based on this insight, we constructed a six-dimensional dynamical model, which analytically predicted cortex-wide mesoscopic activity on single trials from the activity of six driver areas and connectome, without any parameter fitting. Rich patterns of animal movements can be decoded from the activity generated by the model. The six-dimensional model generalized across individual animals and behaviors, indicating that the connectome fundamentally defines the cortex-wide neural dynamics on single trials. Our results reveal that rich spatiotemporal activity across the mouse cortex is confined to the low-dimensional subspace spanned by the driver areas, with the embedding and dynamics determined by the connectome.

3-081. Differential representational drift as a function of stimulus-dependent circuit architecture in V1

Tyler Marks[1,2] TYLERMARKS@UCSB.EDU
Michael Goard[3] MICHAEL.GOARD@LIFESCI.UCSB.EDU

[1]University of California, Santa Barbara
[2]Neuroscience Research Institute
[3]University of California, Santa Babara

For an animal to produce consistent behaviors necessary for survival, the brain must generate stable neural representations of sensory input. However, neural activity in many regions is highly variable, stemming from both trial-to-trial variability and representational drift across days to weeks. Recently, longitudinal imaging studies have begun to address the stability of stimulus representation in sensory cortex, generally finding relatively stable neural responses across experimental sessions in primary sensory areas. However, the majority of chronic recordings in V1 have been performed with artificial visual stimuli, and it is still unclear whether responses to naturalistic stimuli would exhibit similar levels of stability. In this study, we performed chronic 2-photon imaging of populations of neurons in mouse V1 to directly compare the representational stability of artificial versus naturalistic visual stimuli over up to 42 days (Fig. 1A-C). Responses to passive drifting gratings were highly stable across sessions, showing minimal change in tuning across sessions (Fig. 1D). However, in contrast to highly stable artificial stimulus responses, we found that single neuron responses to naturalistic movies exhibited progressive representational drift across sessions, evident even in neurons responsive to both stimuli (Fig. 1D-E). This phenomenon was present across cortical layers as well as in inhibitory interneurons. Control experiments showed that the representational drift in natural movie responses was not due to response magnitude, eye movements, or arousal. This differential stability was accompanied by differential drift in local between-neuron functional connectivity across stimuli (Fig. 1F-H). These results suggest that representational stability in V1 is stimulus-dependent and is driven by the strength of pre-existing circuit architecture of co-tuned neurons (Fig. 1I).

3-082. Stable representation of a naturalistic movie emerges from episodic activity with gain variability

Ji Xia[1,2] XIAJI@WUSTL.EDU
Tyler Marks[3] TYLER.MARKS@LIFESCI.UCSB.EDU
Michael Goard[3] MICHAEL.GOARD@LIFESCI.UCSB.EDU
Ralf Wessel[1] RW@WUSTL.EDU

[1]Washington University In St. Louis
[2]Department of Physics
[3]University of California, Santa Babara

Visual perception of identical visual stimuli remains largely stable throughout our lifetime. Yet, accumulating evidence from longitudinal recordings in cerebral cortex of experimental animals indicates that neural responses to identical sensory stimuli tend to vary across trials and can vary even more across experimental sessions. How stable internal representations of sensory stimuli arise from unstable neural responses is presently unclear. To address this question, we performed longitudinal two-photon calcium imaging of excitatory neurons in V1 of awake, head-fixed mice during visual stimulation with repeated identical naturalistic movie clips across weeks. Single neuronal responses consisted of sparse episodic activity which were stable in spike time but unstable in spike rates across weeks (Fig. A, B). By fitting a linear model to capture the population activity, we found that the individual spiking episode, instead of neuron, served as the basic unit of the week-to-week fluctuation. The spiking episodes evenly tiled the duration of the natural movie, each following a distinct gain change pattern across weeks, even for episodes within the same neuron. To investigate how the high-dimensional population activity with fluctuations encodes the stimulus, we directly visualized the neural manifold (Fig. C). Using an unsupervised method, we revealed a stable one-dimensional representation of the time in the natural movie (Fig. D). Moreover, the majority of the week-to-week fluctuation was perpendicular to the stimulus encoding direction, thus leaving the stimulus representation largely unaffected. By analyzing surrogate neuronal population activity with preserved trial structure, we found that the precise timing of episodic activity was essential for constraining the direction of neural fluctuation, while coordination among episodic activity was key to uniquely representing time points in the trial. We propose that precise episodic activity with coordinated gain changes are keys to maintain a stable stimulus representation in V1, despite unstable single neuron activity.

3-083. Decentralized motion inference and registration of Neuropixels data

Erdem Varol[1] EV2430@COLUMBIA.EDU
Julien Boussard[1] JB4365@COLUMBIA.EDU
Nishchal Dethe[1] ND2506@COLUMBIA.EDU
Liam Paninski[2] LIAM@STAT.COLUMBIA.EDU

[1]Zuckerman Mind Brain Behavior Institute, Columbia University
[2]Columbia University

Multi-electrode arrays are essential tools for studying the spiking activity of individual neurons across the brain. Neuropixels, dense linear probes with hundreds of regularly spaced electrodes, have recently pushed the state of the art forward significantly. During recordings, Neuropixels probes shift in position due to subtle movements experienced by the brain. Inferring the amount of movement that is experienced by the Neuropixels array and correcting the voltage signal to account for this movement is a crucial step towards increasing the stability and signal to noise ratios for any downstream analyses. Previous work has addressed the difficulty of registering data with sparse neural firing patterns by representing the signal as a spatial histogram of voltage peak-to-peak values patterns broken into time bins; by collating the signal across a chunk of time into a histogram, we obtain a data representation in which the neurons' spatial location and intensity are the salient features used to perform registration. Here, to address the variable firing patterns in subpopulations of neurons, we introduce the notion of decentralized registration. Unlike traditional registration techniques, we treat each frame as its anchor and estimate relative displacements with respect to every pair of frames. We find that this novel decentralized registration framework improves displacement estimation and signal interpolation in synthetic and experimental datasets compared to template-based registration and state-of-the-art methods.

3-084. Sequential neurogenesis shapes hippocampal network dynamics

Roman Huszar[1] ROMAN.HUSZAR.17@GMAIL.COM
Yunchang Zhang[1] YZ5230@NYU.EDU
Heike Blockus[2] HB2505@COLUMBIA.EDU
Gyorgy Buzsaki[1] GYORGY.BUZSAKI@NYULANGONE.ORG
[1]New York University
[2]Columbia University

The hippocampus plays a key role in the storage and recall of episodic memories. However, the mechanisms by which these functions occur are a subject of debate. According to one hypothesis, hippocampal networks incorporate new information through activity-dependent plasticity. An alternative view suggests that developing networks give rise to a reservoir of activity patterns that are available for matching with experiences. The origin and organization of such preexisting activity patterns remain unexplored. Here, we show that hippocampal network dynamics are preconfigured by sequential neurogenesis. We labeled pyramidal neurons according to their embryonic birthdate by transfecting hippocampal progenitors with channelrhodopsin. Neurons spanning a wide range of embryonic birthdates (E13, E14, E15, and E16) were identified optogenetically in adult mice implanted with high density silicon probes. Pairwise correlations between neurons born at the same time were stronger than between neurons born at different dates. Furthermore, cell assembly patterns, defined by repeated neural coactivations, were expressed at different rates depending on the birthdate of assembly members. These differences tracked nonlinear changes in neurogenesis, in that higher neurogenesis rates were associated with stronger assembly expression. The rate of assembly expression was predictive of spatial receptive field tuning during behavior. In particular, strongly expressed assemblies tended to form a larger number of reliable place fields rather than single fields. At the single neuron level, these differences reflected the number of place fields of birthdated neurons. We hypothesize that nonlinear neurogenesis rates promote the formation of assemblies with distinct receptive field tuning. A neuron's propensity for forming place fields may be biased by its membership in preconfigured assembly patterns.

3-085. A Bayesian framework for generalized entrainment to stochastic rhythms

Jonathan Cannon[1,2] JCAN@MIT.EDU
[1]Massachusetts Institute of Technology
[2]Brain and Cognitive Science

When presented with complex rhythmic auditory stimuli, most humans are able to track underlying temporal structure (e.g., a "beat") with overt movement or covertly while automatically and unconsciously adjusting for timing irregularities. We propose that the problem of overt and covert rhythm tracking is most naturally characterized as a problem of point process filtering: continuously estimating a hidden underlying phase and tempo based on information provided by the precise timing of temporally localized events. Event timing is informative because events are expected to occur at certain phases with certain probabilities and degrees of temporal precision, as specified by a flexible periodic or aperiodic "temporal expectation template." We demonstrate that approximate (variational Bayesian) solutions to this inference problem reproduce characteristics of overt and covert human rhythm tracking that have not been well addressed by other models, including interval-dependent phase correction, failure to track overly syncopated rhythms, and the distortion of perceived timing by disappointed expectations.

This framework can treat rhythms of arbitrary complexity, and can therefore be used model the behavioral results of rich musical psychophysics experiments. Further, its continuous time dynamics serve as a plausible model of brain activity during entrainment. This work is motivated by a neurophysiologically detailed hypotheses of the brain dynamics of entrainment, and a search for brain activity similar to these inference dynamics could build a link between the algorithmic and physiological levels of entrainment. Finally, as a variational Bayesian inference method in continuous time, this approach fits comfortably into the "predictive processing" framework, allowing for future extension into models of hierarchical inference based on event timing that could describe even more complex sensorimotor processes.

3-086. NP-GLM: Nonparametric GLM

Matthew Dowling[1,2] MATTHEW.DOWLING@STONYBROOK.EDU
Yuan Zhao[1] YUAN.ZHAO@STONYBROOK.EDU
Memming Park[1,2] MEMMING.PARK@STONYBROOK.EDU
[1]Stony Brook University
[2]Neurobiology and Behavior

GLMs are a prevalent way to model neural encoding. However, for the practitioner correctly parameterizing and fitting these models can often be an uphill battle. This is because the standard approach when fitting these models to neural spiking data is to use a set of temporal bases functions for each variable of interest. At minimum this involves specifying a particular functional form of each basis, a temporal window, and the number of bases used. Misspecification of these parameters, foregoing extensive model selection, or hand tuning the parameters can lead to poor inference and misguided scientific conclusions. Often, heavy regularization will also be a necessity especially when examining neurons with low firing rates. To circumvent the difficulties that arise in the aforementioned procedure we propose a nonparametric approach for jointly learning the filters and their hyperparameters by using Gaussian Process priors – Nonparametric GLM (NP-GLM). Taking advantage of the sparse variational inference framework, through the use of a set of inducing points, our proposed method remains computationally efficient while still providing flexible and expressive characterization of GLM filters. Moreover, optimization of inducing point locations endows us with a way of automatically learning the appropriate temporal span of each individual filter. This culminates in the proposed method requiring less specification and being more apt to "plug-and-play" data analysis. While Bayesian methods have been used before in the same vain to model neural encoding, they have also characterized the filters using a fixed set of bases, not alleviating the difficulties discussed. To demonstrate the efficacy of our method we analyze synthetic data as well as RGC responses and show it outperforms the standard approach in GLM fitting while remaining robust to any parameterizations

3-087. A distributed circuit for associating environmental setting to motor choice in retrosplenial cortex

Luis Franco[1,2] LUIS.FRANCO@LIFESCI.UCSB.EDU
Michael Goard[3] MICHAEL.GOARD@LIFESCI.UCSB.EDU

[1]University of California, Santa Barbara
[2]Neuroscience Research Institute
[3]University of California, Santa Babara

During navigation, animals often use recognition of familiar environmental settings to guide motor action selection. The retrosplenial cortex (RSC) receives inputs from both visual cortex and subcortical regions required for spatial memory, and projects to motor planning regions. However, it is not known whether RSC is important for associating familiar environmental settings with specific motor actions. Here, we test this possibility by developing a task in which trajectories are chosen based on the current environmental setting. We find that mice exhibit differential pre-decision activity in RSC, and that optogenetic suppression of RSC activity impairs task performance. Indeed, individual RSC neurons encode a range of task variables, often multiplexed with distinct temporal profiles. Moreover, their responses are spatiotemporally organized, with different task variables represented along specific posterior-to-anterior gradients during behavioral performance. These results reveal an anatomically-organized RSC circuit for associating environmental settings to appropriate motor outputs during decisions.

3-088. Towards a large-scale model of sleeping brain: sleep spindles in network mass models

Nikola Jajcay[1,2] JAJCAY@NI.TU-BERLIN.DE
Klaus Obermayer[1] KLAUS.OBERMAYER@TU-BERLIN.DE

[1]Technische Universitat Berlin
[2]Neural Information Processing Group

The hierarchical nesting of sleep spindles and slow oscillations is considered a precursor of successful episodic memory consolidation where it, presumably, sets the stage for memory traces migration from short-term hippocampal storage to longer-lasting neocortical sites[1]. Here we take the first steps towards integrating thalamocortical projections into a large-scale brain model and modelling whole-brain cortical slow-wave activity and thalamic spindles as seen in non-REM sleep. We model the thalamus as a node containing one excitatory and one inhibitory mass representing thalamocortical relay neurons and thalamic reticular nuclei, respectively[2]. We found dynamically interesting spindle-promoting regimes and we were also able to control the inter-spindle interval, spindle duration, and shape of spindle envelope. Secondly, we connected the thalamic node to one cortical node, modelled as excitatory and inhibitory adaptive exponential integrate-and-fire neuronal masses[3]. The excitatory mass contains a spike-triggered adaptation mechanism and is parametrised to sit in the limit cycle generating slow oscillations. Moreover, the cortical node is adapted to include perturbation via transcranial current stimulation. Our results when investigating the dynamical repertoire show that, upon connection, the thalamic spindles are imprinted into cortical node activity where they modulate slow oscillation envelope, and that slow os-

cillation activity of the cortex, in turn, shapes spindling behaviour of thalamocortical relay mass by affecting spindle duration and inter-spindle interval. Our connected model also conserves the phase-phase and phase-amplitude couplings reported in the literature on observed EEG data. Our results suggest that thalamic mass model of spindle activity can be connected to various mass or mean-field models of cortical nodes and, after careful treatment of network connections and delays, we believe that our conclusions would carry over to the large-scale network model.

3-089. Cortical wide-field imaging reveals the dominant role of hippocampal Gamma-events in orchestrating modular activation

Federico Stella F.STELLA@DONDERS.RU.NL
Francesco P. Battaglia FPBATTAGLIA@GMAIL.COM
Rafael Pedrosa RHAPEDROSA@GMAIL.COM

Donders Institute for Brain, Cognition and Behaviour, Radboud University

Recent advancements in anatomical and functional imaging techniques have led to a profound reappraisal of the principles of cortical mapping, emphasizing whole-brain networks that are organized according to functional and connectivity gradients.

Yet, how spontaneous activity propagates in the neocortex remains to be understood. Moreover, hippocampal-cortical interactions are thought to be crucial for memory processes, but are typically studied with focus on a handful of putative “hubs”. Here, we investigate the fine-grained spatiotemporal dynamics of spontaneous activity in the entire dorsal cortex with simultaneous recordings of wide-field voltage sensitive dye transients (VS), cortical ECoG, and hippocampal LFP in anesthetized mice.

Both VS and ECoG show cortical avalanches. When measuring avalanches from the VS signal, we find a major deviation of the size scaling from the power-law distribution predicted by the criticality hypothesis and well approximated by the results from the ECoG. Breaking from scale-invariance, avalanches can thus be grouped in two regimes. Small avalanches consists of a limited number of co-activation modes involving a sub-set of cortical networks (related to the Default Mode Network), while larger avalanches involve a substantial portion of the cortical surface and can be clustered into two families: one immediately preceded by Retrosplenial Cortex activation and mostly involving medial-posterior networks, the other initiated by Somatosensory Cortex and extending preferentially along the lateral-anterior region.

Rather than only differing in terms of size, these two set of events appear to be associated with markedly different brain-wide dynamical states: they are accompanied by a shift in the hippocampal LFP, from the ripple band (smaller) to the gamma band (larger avalanches), and correspond to opposite directionality in the cortex-to-hippocampus causal relationship. These results provide a concrete description of global cortical dynamics, and shows how cortex in its entirety is involved in bi-directional communication in the hippocampus even in sleep-like states.

3-090. An unified neural circuit model for travelling wave, phase precession, and anticipative tracking

Tianhao Chu[1,2] 2001111555@STU.PKU.EDU.CN
Si Wu[1] SIWU@PKU.EDU.CN
Yuanyuan Mi[3] MIYUANYUAN0102@CQU.EDU.CN

[1]Peking University
[2]Academy for Advanced Interdisciplinary Studies
[3]Chongqing University

Continuous attractor neural networks (CANNs) have been widely used as a canonical model for neural information processing. Travelling wave in the cortex, phase precession of hippocampus place cells at the theta rhythm, and anticipative tracking of head direction were observed in neurophysiological experiments. Previously, the three phenomena were modelled separately and their link was not recognized. Here, by including negative feedback in the dynamics of a CANN, we show that the three phenomena can be unified in a single computational framework under different parameter regimes. We consider spike frequency adaptation (SFA) as an example of negative feedback. By fixing the amplitude of SFA, we find that the CANN exhibits three dynamical behaviors in response to external moving inputs, which are: 1) when the external input is very small, the network holds travelling wave, in term of that the neural bump (localized neural activities) propagates spontaneously in the space; 2) when the external input is very strong, the neural bump tracks the input anticipatively, in term of that the bump position leads the input location by a constant time, reminiscing anticipative tracking for head direction; 3) when the external input is of appropriate size, the neural bump tracks the input in an oscillatory manner, in term of the bump position oscillates around the input location. If we inspect the activity of an individual neuron during oscillatory tracking, the phase precession phenomenon similar to that of hippocampus place cells is observed. Overall, our study reveals that the seemingly unrelated neural behaviors of travelling wave, phase precession, and anticipative tracking can be unified in a single CANN with negative feedback under different parameter regimes. We hope that this study will give us insight into understanding the general mechanism of neural information processing.

3-091. Memory improvement by consolidation via synaptic tagging and capture in recurrent neural networks

Jannik Luboeinski[1,2] JANNIK.LUBOEINSKI@PHYS.UNI-GOETTINGEN.DE
Christian Tetzlaff[1] TETZLAFF@PHYS.UNI-GOETTINGEN.DE

[1]University of Goettingen
[2]Departmental of Computational Neuroscience

The synaptic-tagging-and-capture (STC) hypothesis formulates that at each synapse the concurrence of a tag with protein synthesis yields the maintenance of changes induced by synaptic plasticity. This hypothesis provides a biological principle underlying the synaptic consolidation of memories that is not verified for recurrent neural circuits. We developed a theoretical model integrating the mechanisms underlying the STC hypothesis with calcium-based synaptic plasticity in a recurrent spiking neural network. In the model, calcium-based synaptic plasticity yields the formation of strongly interconnected cell assemblies encoding memories, followed by consolidation through the STC mechanisms. To evaluate the functionality of a memory, we measure the recall quality. To this end, we employ two different means: the mutual information between the activity distributions during recall and learning, and a coefficient describing the completion of an input-defined pattern. Using these two quantities, we explore the memory functionality varying the strength of inhibitory connections and the size of the Hebbian cell assembly, and thereby show the capability of STC to consolidate memories in recurrent neural networks. In addition, we show for the first time that STC mechanisms lead to significant improvement of memory recall after several hours by modifying the storage of memories with the passage of time. Examining the underpinnings of this effect, we find that it is in fact caused by two processes that act together: by a passive improvement during consolidation and by an active improvement during recall of the consolidated memory. This kind of memory enhancement can provide a new principle for storing information in biological and artificial neural circuits, and could be linked to pertinent existing psychological findings.

3-092. The generation of cortical novelty responses through inhibitory plasticity

Christoph Miehl[1] CHRISTOPH.MIEHL@BRAIN.MPG.DE
Auguste Schulz[2] AUGUSTE.SCHULZ@UNI-TUEBINGEN.DE
Michael Berry II[3] BERRY@PRINCETON.EDU
Julijana Gjorgjieva[1,4] GJORGJIEVA@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]University of Tuebingen
[3]Princeton Neuroscience Institute
[4]Computation in Neural Circuits Group

The ability to predict the near and far future is crucial for the survival of an animal. The brain is known to perform predictive computations upon incoming sensory stimuli. How these predictive computations are implemented mechanistically, however, remains unknown. One way to test for predictive computations is to create violations of periodic stimulus sequences. Several experiments have shown that novel or rare stimuli elicit strong responses in an otherwise adapted neural response. Here, we explore the potential of inhibitory synaptic plasticity as a biologically-plausible mechanism for the generation of adapted and novelty responses. Recent experimental studies point towards an essential role of inhibition and plasticity between inhibitory and excitatory neurons, but their contribution has been underexplored so far. By implementing a sequence violation paradigm in a computational network model of spiking neurons, we show that inhibitory plasticity is sufficient to capture novelty responses and adaptive phenomena on multiple timescales. We show that generated novelty responses do not depend on the exact temporal structure of a sequence but rather on the distribution of presented stimuli. In contrast to most previous plastic recurrent network models, we implement tuning of both the excitatory and inhibitory population. Through the addition of stimulus-tuning of inhibitory cells, the model further captures stimulus-specific adaptation (SSA). This finding reveals that the network configuration encompasses computational capabilities beyond those of intrinsic adaptation. Furthermore, we suggest disinhibition to be a powerful regulator, controlling the amplifi-

cation of novelty responses. Our model can capture a large range of experimental data concerned with novelty detection. In summary, we demonstrate how inhibitory plasticity, a mechanism commonly introduced to stabilize plastic recurrent neural networks, can give rise to multiple interesting circuit computations.

3-093. Combining computational controls with natural text reveals new aspects of meaning composition

Mariya Toneva[1,2] MKTONEVA@CS.CMU.EDU
Tom Mitchell[1] TOM.MITCHELL@CS.CMU.EDU
Leila Wehbe[1] LWEHBE@CMU.EDU

[1]Carnegie Mellon University
[2]Machine Learning, Neural Computation

Understanding language in the real-world requires us to compose the meaning of individual words in a way that makes the final composed product more meaningful than the string of isolated words. For example, we understand the statement that “Mary finished the apple" to mean that Mary finished eating the apple, even though “eating" is not explicitly specified. This supra-word meaning is at the core of language comprehension, and its neurobiological bases and processing mechanisms must be specified in the pursuit of a complete theory of language processing in the brain. The current work provides an operational definition of supra-word meaning by defining it as the multi-word meaning that is beyond the meaning of individual words. We further provide a computational representation of supra-word meaning using powerful neural network algorithms and a new approach to control for shared information between supra- and individual-word meaning. Using fMRI recordings, we reveal that hubs thought to process lexical-level meaning also maintain supra-word meaning, suggesting a common substrate for lexical and combinatorial semantics. However, surprisingly, we find that supra-word meaning is not detectable in MEG. The difference between the fMRI and MEG results suggests that the processing of supra-word meaning is based on neural mechanisms that are not related to synchronized cell firing, as is the MEG signal. We show that what MEG recordings are sensitive to instead is information that is unique to the individual recently-read words. These results call for a more nuanced understanding of previous works that aim to study composition of sentence-level meaning using imaging modalities that rely on synchronized firing (i.e. MEG, EEG, ECoG). This finding also has important implications for building brain computer interfaces (BCI) to decode meaning from brain recordings in the real world.

3-094. Does Nature play dice? A comparative study of two Drosophila olfactomes

Hamza Giaffar[1,2] HGIAFFAR@CSHL.EDU
Sergey Shuvaev[3] SSHUVAEV@CSHL.EDU
Dmitry Rinberg[4] RINBERG@NYU.EDU
Alexei Koulakov[1] AKULA@CSHL.EDU

[1]CSHL
[2]Neuro
[3]Cold Spring Harbor Laboratory
[4]New York University

The peripheral olfactory system analyses odorous chemical mixtures into activity in a large number of genetically defined information channels, each associated to an olfactory receptor (OR) type. These channels are combined two synapses downstream in putatively associative areas, including the mammalian piriform cortex and insect mushroom body. The logic of this combination is an important algorithmic feature of the olfactory circuit and will be the focus of this work. According to the dominant view in the field, disordered or random combinations are an important computational ingredient in olfactory information processing; in this view, randomness was discovered and selected for during the evolution of olfaction. Another view argues that there may be some non-trivial organization in the combination of OR types, perhaps encoding information about environmental statistics accumulated over the course of evolution, or another algorithmic trick to better distill the complex chemical world into appropriate behaviour. These opposing views can, and to a limited extent have been, tested on data. Here we advance this effort by comparing two recently reported and largely complete D. Melanogaster connectomes, finding, in contrast to earlier work, significant structure or non-randomness that is consistent between datasets. We describe this shared non-random component of connectivity and discuss its possible relation to olfactory primacy coding. Finally, we discuss attempts to understand these results in computational terms, i.e. we address the question what is it good for?

3-095. Nonlinear intensity modulation allows biologically plausible independent component analysis

Yanis Bahroun[1] YBAHROUN@FLATIRONINSTITUTE.ORG
Anirvan M. Sengupta[2] ANIRVANS.PHYSICS@GMAIL.COM
Dmitri B. Chklovskii[3] MITYA@FLATIRONINSTITUTE.ORG

[1]Flatiron Institute, Simons Foundation
[2]Rutgers University, Flatiron Institute
[3]Flatiron Institute, Simons Foundation, NYU Neuroscience Institute

Various brain areas, particularly the early auditory, olfactory, and visual systems, are known to effortlessly solve the cocktail party problem. This problem is often addressed by Independent Component Analysis (ICA). The development of biologically plausible neural circuits capable of solving ICA may shed light on general neural computation principles of unsupervised extraction of independent sources from mixtures. We search for a solution that can be implemented using neural networks (NN) derived from a fully normative approach, i.e., both the neural architecture and synaptic rules must be derived from an optimization problem. The NN should also operate in the online setting where the dataset is streamed and be trained with local learning rules to be biologically plausible. This approach has been achieved in the limited settings of nonnegative ICA and bounded component analysis. We introduce a novel similarity matching (SM) objective for ICA from which we derive a two-layer NN. The synaptic weights in our networks are updated using local learning rules. Interestingly, the learning rules in the second layer are nonlinearly modulated by the overall intensity of the layer's inputs. Also, the resulting NN is reminiscent of early sensory processing stages in the brain.

3-096. Neural circuits for distinct locomotor behaviors in Drosophila larvae

Marina Wosniack[1] MARINA-ELAINE.WOSNIACK@BRAIN.MPG.DE
Zhuoshi Liu[2] ZHUOSHI.LIU@UNI-OLDENBURG.DE
Aishwarya Patel[3] AISHWARYA.PATEL@TUM.DE
Konrad Heinz[4,5] KONRAD_HEINZ@HOTMAIL.DE
Jimena Berni[6] J.BERNI@SUSSEX.AC.UK
Julijana Gjorgjieva[1,7] GJORGJIEVA@BRAIN.MPG.DE

[1]Max Planck Institute for Brain Research
[2]Carl von Ossietzky University of Oldenburg, Max Planck Institute for Brain Research
[3]Technical University of Munich, Max Planck Institute for Brain Research
[4]University of Exeter
[5]University of Cambridge
[6]Brighton and Sussex Medical School, University of Sussex, University of Cambridge
[7]Computation in Neural Circuits Group

Drosophila larvae perform sequences of crawls, turns, and pauses to navigate in the substrate. The movement is generated by muscle contraction waves that propagate along the larval body. Symmetric contraction of the left and right sides of the body results in crawls, while the asymmetric contraction of the most anterior segments generates a turn. The isolated central nervous system of the larva generates spontaneous activity waves that alternate between symmetric and asymmetric propagation – the same motifs found in crawls and turns. Coordinated waves have been observed even when the brain was acutely silenced, suggesting that central pattern generators (CPGs) for crawls and turns are distributed in the ventral nerve cord (VNC). Previous studies addressed the generation and propagation of symmetric activity waves, while the mechanisms behind the generation of asymmetric waves remain elusive.

We combined a detailed analysis of calcium imaging recordings of spontaneous waves in the larval VNC and modeling to investigate (1) properties of symmetric and asymmetric waves and (2) connectivity motifs in a network model that can mediate the generation of both types of waves. We found that periods of asymmetric activity usually follow or are followed by symmetric backward waves. When comparing symmetric forward and backward waves we found that the relative duration of the activation of posterior and anterior segments is different for each type of wave. These observations combined with the recent literature inspired our network model based on CPGs with different connectivity motifs. We test the effect of those motifs in the properties of the generated symmetric and asymmetric waves and make hypotheses about the organization of these neurons in the larval VNC. Beyond generating spontaneous locomotor waves, our model can be used to study goal-directed behavior by including descending neurons that modulate the dynamics of motor patterns in the VNC.

3-097. Distinct directional modulations of hippocampal theta scale spike timing in areas CA1 and CA3

Yuk Hoi Yiu[1,2] HOI.YIU@BIO.LMU.DE
Jill Leutgeb[3] JLEUTGEB@UCSD.EDU
Christian Leibold[1] LEIBOLD@BIO.LMU.DE

[1]Ludwig Maximilian University of Munich
[2]Computational Neuroscience, Department Biology II
[3]Neurobiology Section and Center for Neural Circuits and Behavior, University of California, San Diego

Place cells in the rodent hippocampus are able to represent distance information by a temporal code. They fire at progressively earlier phases with respect to the LFP theta oscillation (4-12Hz) as the animal traverses the place field. This phenomenon is known as phase precession and is thought to partially underlie pair-wise correlation structure and, hence, theta sequences. Apart from position, some place cells also encode heading direction by modulations of place field firing rate. Although rigid firing sequences should necessarily induce directionality, it has not been studied so far whether phase precession and pair-wise correlations are also modulated by directional input.

Here, we investigated this question by re-analyzing previously published spiking data from hippocampal CA1, CA2 and CA3 pyramidal neurons in rats performing a 2D free foraging task (Mankin et al., 2015). Our results indicate that phase precession preferentially occurs in a direction opposing the best firing rate direction by 180 degree. Phase precession is most directional in CA3 and the best precession direction is associated with high onset phases (i.e. more prospective firing) and a shallower slope. Furthermore, we found that theta-scale correlation lags in CA3 are less influenced by directional input as compared to CA1, showing more reliance on inherent network dynamics in agreement with the recurrent connections in CA3. However, simulating place cell activity using the model by Romani and Tsodyks (2015) and incorporating the measured animal trajectories could not reproduce the experimental results, which indicates that the directional modulation of phase precession and correlations in CA3 cannot be solely accounted for by the asymmetrical weight couplings between place cells. In summary, our findings reveal that there is an interaction between directionality and theta scale correlations. Further extension of current hippocampal models is required to account for such a relationship.

3-098. Task contingencies determine the temporal window of causal involvement of mouse V1 in perception

Matthijs oude Lohuis[1,2] M.N.OUDELOHUIS@UVA.NL
Jean Pie[1] J.L.PIE@UVA.NL
Pietro Marchesi[1] P.MARCHESI@UVA.NL
Christiaan de Kock[3] CKK230@VU.NL
Cyriel Pennartz[1] C.M.A.PENNARTZ@UVA.NL
Umberto Olcese[1] U.OLCESE@UVA.NL

[1]University of Amsterdam
[2]Swammerdam Institute for Life Sciences
[3]Vrije Universiteit Amsterdam

In primary sensory cortices, stimulus presentation triggers an early response component (driven by thalamic bottom-up input) and, often, a late, outcome-related component, thought to result from reverberating activity through areal interactions. Across primates and rodents, this late activity correlates with reports of perception. However, a causal link with behavioral outcome is only found in tactile but not visual perception (1–3). Here we wondered whether the presence and causal contribution of this late activity depends on task context. We investigated how the same visual stimulus (moving gratings at a given orientation) is processed as a function of task-related context: the saliency of a stimulus, whether it is predictive of reward, and whether it is embedded in a multisensory task context. We trained three cohorts of mice in an audiovisual change detection paradigm with same stimulus configurations, but different reward contingencies: naive (untrained), unisensory trained (detecting visual, but ignoring auditory), or multisensory trained (detecting both visual and auditory). While visual hit rates were similar for both trained cohorts , multisensory-trained mice responded more slowly to the same visual stimuli than unisensory-trained mice. Laminar recordings in V1 revealed that early-latency, sensory-related activity was similar across cohorts, whereas late, outcome-related activity only emerged in trained mice and consistently preceded the first lick by 280 ms. To test whether this late activity was causal to perception, we optogenetically silenced V1 bilaterally either at stimulus onset (0 ms) or after a delay (200 ms). Early silencing indiscriminately abolished perceptual report, but delayed silencing selectively affected sessions in multisensory-trained mice with slow reaction times and a late onset of outcome-related activity in V1. A multisensory task context thus delays reaction time and decision-related activity in V1 and this determines whether late activity in V1 only correlates or is necessary to perceive and respond to visual stimuli.

3-099. Self-organization of discrete grid cell modules from continuous gradients in connectivity

Mikail Khona[1,2] MIKAIL@MIT.EDU
Sarthak Chandra[1] SARTHAKC@MIT.EDU
Francisco Acosta[1,3] FACOSTA@MIT.EDU
Ila R. Fiete[1] FIETE@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Brain and Cognitive Sciences, Physics
[3]Department of Physics

Grid cells of the mammalian medial entorhinal cortex (mEC) exhibit periodic lattice-like tuning curves in their encoding of space as animals navigate the world. Nearby grid cells have identical lattice periods, but at larger separations along the long axis of mEC the period jumps in discrete steps so that the full set of periods clusters into 5-7 groups called modules. These modules endow the grid code with many striking properties such as an exponential capacity to represent space and an unprecedented robustness to noise. However, the formation of discrete modules is puzzling given that the biophysical properties of mEC stellate cells – including inhibitory inputs from PV interneurons, time constants of EPSPs and spike after-hyperpolarization, intrinsic resonance frequency and the propensity to burst and also differences in gene expression – vary in continuous topographic gradients along the long axis. How does discreteness in grid modules arise from continuous gradients? We study this question using a continuous attractor network model, the generally accepted hypothesis for the working of grid cells. In these networks, a Turing pattern formation process leads to creation of hexagonal firing patterns. We show that the addition of a continuous gradient in local inhibition along the long axis, together with a weaker second local fixed-range inhibitory interaction, drives modularity in the pattern formation process such that the lattice cluster into a few discrete periods. We show that the emergence of modularity is tied to the emergence of localized eigenmodes that interact with the delocalized lattice formation eigenmodes. A one dimensional Kuramoto model with local coupling and a continuous gradient in natural frequency provides a simplified model of modularity emergence.

3-100. Neurons, astrocytes, and anything in between

Erik Peterson[1,2] ERIK.EXISTS@GMAIL.COM
Tim Verstynen[1] TIMOTHYV@ANDREW.CMU.EDU

[1]Carnegie Mellon University
[2]Psychology

Why should artificial networks use neurons with complex morphologies — why does nature? In this work we'll first outline a new differential continuum of cellular specialization consisting of compartments, processes, waves/spikes, and a generalization of synapses. We meld biological and artificial learning, in other words. We want to move smoothly from networks of point cells, who communicate using waves of activity (like astrocytes), to extremely complex single neurons (akin to pyramidal cells), and anything in between. Second, we'll describe some new preliminary simulations, where we aim to uncover basic tradeoffs between cellular specialization, speed of learning, and generalization performance. Our testbed is a large set of simple randomly generated toy tasks.

3-101. A nonlinear calcium-based plasticity rule explains long-term stability of synapses

Bin Wang[1,2] BIWANG@UCSD.EDU
Johnatan Aljadeff[1,3] ALJADEFF@UCSD.EDU

[1]University of California, San Diego
[2]Physics
[3]Division of Biological Sciences, Section of Neurobiology

Activity dependent synaptic plasticity is considered to be a primary mechanism underlying learning and memory. Yet it is unclear whether plasticity rules such as STDP measured in vitro apply to in vivo conditions. Network models with STDP predict that activity patterns (e.g., place-cell spatial selectivity) should change much faster than observed experimentally. Here we address this gap by investigating a nonlinear calcium-based plasticity rule, which was fit to experiments done in physiological conditions. In this model, long-term potentiation and depression result from intracellular calcium transients. Crucially, elevated calcium transients arise almost exclusively from

synchronous coactivation of pre- and postsynaptic neurons. We analytically approximate the full distribution of nonlinear calcium transients as a function of pre- and postsynaptic firing rates, and temporal correlations. This analysis directly relates activity statistics that can be measured in vivo to the changes in synaptic efficacy they cause. In contrast to previous models of STDP, our results suggest that both high-firing rates and strong temporal correlations can lead to significant changes to synaptic efficacy. The model's analytical tractability allows us to explore its implications on the network level. Using a mean-field approximation, we show that the nonlinear plasticity rule, without any fine-tuning, gives a stable, unimodal synaptic weight distribution characterized by a large fraction of strong synapses, as seen in experiments. Individual synapses within this distribution remain stable over long periods of time, consistent with previous behavioral studies. Our results provide a mechanistic understanding of how stable learning and memory emerge on the behavioral level from an STDP rule measured in physiological conditions. Furthermore, the plasticity rule we investigate is mathematically equivalent to other learning rules which rely on the statistics of coincidences, so we expect that our formalism will be useful to study other learning processes beyond the calcium-based plasticity rule.

3-102. Decomposition of higher-order correlations by nonlinear Hebbian learning

Gabriel Ocker[1,2] GKOCKER@BU.EDU
Michael Buice[3] MICHAELBU@ALLENINSTITUTE.ORG

[1]Boston University
[2]Mathematics and Statistics
[3]Allen Institute

It is a truth universally acknowledged that a linear neuron in possession of synaptic scaling and a Hebbian learning rule on its inputs must be in want of their first principle component [Oja 1982, Oja and Karhunen 1985]. This observation jump-started an explosion of work in unsupervised feature learning by neural networks and is a hallmark example of a neural computation. Biological synaptic plasticity, however, exhibits nonlinearities that are not accounted for by classic Hebbian learning rules. The computational impact of these nonlinearities has been studied primarily by simulations and heuristic arguments. Here, we introduce a simple nonlinear generalization of Hebbian learning and study it analytically. To focus on the computation implemented by these dynamics, we study the simplest setting of a neuron receiving feedforward inputs.

We show that this nonlinear Hebbian rule allows a neuron to learn tensor decompositions of its higher-order input correlations. The particular input correlation decomposed and the form of the decomposition depend on the location of nonlinearities in the plasticity rule. For simple, biologically motivated parameters steady states of the plasticity correspond to tensor eigenvectors of higher-order input correlations. We prove that the first tensor eigenvector is a globally attracting fixed point of the nonlinear Hebbian dynamics, thus showing that nonlinear Hebbian learning allows a neuron to encode higher-order input correlations in a simple fashion. We then study the impact of fluctuations in the neural activity and in the synaptic weights on the plasticity dynamics by developing a path integral formulation for the joint density functional of the activity and synaptic weights.

3-103. Functional circuitry for goal-directed movements in the superior colliculus

Jaclyn Essig[1,2] JACLYN.ESSIG@CUANSCHUTZ.EDU
Gidon Felsen[1] GIDON.FELSEN@CUANSCHUTZ.EDU

[1]University of Colorado- Anschutz Medical Campus
[2]Physiology and Biophysics

The processes of action selection and execution have traditionally been the focus of separate studies despite their shared neural circuitry. For example, the superior colliculus (SC) has long been known to mediate both the selection of spatial targets and the initiation of movements to acquire these targets; however, the relationship between these processes is unclear. We hypothesized that SC functional circuitry allows target selection to produce a “default template” for the ensuing motor plan such that once a target is selected, a pre-determined pattern of activity mediates target acquisition. We tested this hypothesis by recording from SC neurons, including identified GABAergic neurons, in mice performing a spatial choice task. The SC encodes for specific regions of contralateral space such that more eccentric targets are initially represented more caudally in the SC, with rostral activity emerging as the target is acquired. Our task required mice to select an eccentric left or right reward port based on a binary odor mixture. Importantly, after odor delivery, mice wait for a “go signal” before orienting to the reward port, temporally separating neural activity during target selection from movement-related activity as the target is acquired. Single-unit recordings along the rostrocaudal axis of the SC revealed a large population of neurons that are active during both epochs, with increased rostral activity later in movement, consistent with our hypothesis that SC activity underlying target acquisition is in part determined by activity underlying target selection. We next examined how SC circuitry could link these processes using a biologically plausible attractor model and uncovered intrinsic SC mechanisms that promote a rostral shift in activity during movement. Our work reveals a neural mechanism for how choices and actions can be dynamically encoded by the same neurons in the same circuit.

3-104. Framework to generate more realistic GLM spike trains

Diego Arribas[1,2] DIEGOARRI91@GMAIL.COM
Yuan Zhao[1] YUAN.ZHAO@STONYBROOK.EDU
Memming Park[1,2] MEMMING.PARK@STONYBROOK.EDU

[1]Stony Brook University
[2]Neurobiology and Behavior

To characterize spatiotemporal structure in neural time series, we often use autoregressive models—statistical models that predict the next time step given a window of past observations. This category includes the popular generalized linear model, used for spike trains to determine both the stimulus encoding and dependence among neurons. The standard approach to fitting an autoregressive spike train model is to maximize the likelihood for one-step prediction. However, the maximum likelihood estimation often leads to models that generate spike trains that do not resemble the data on which they were trained on—the model-generated spike trains can fail to capture important features of the data and even show diverging firing rates. This harms their neuroscientific interpretation and it is a limiting factor if model samples are required for subsequent analysis. To alleviate this, we propose a new stochastic fitting procedure that directly minimizes a kernel induced divergence between the observed and model-generated spike trains. We validated our approach on both real and synthetic neural data and obtained well-behaved models. Combinations of spike train kernels enable us to control the trade-off between different features which is critical for dealing with model-mismatch. Our framework is general and can be applied to any autoregressive generative model in neuroscience.

3-105. Gain variability, stimulus uncertainty, and divisive normalization: a normative approach

Jeff Beck[1,2] BRAINS4LUNCH@GMAIL.COM

[1]Duke University
[2]Neurobiology

Behavioral evidence consistently shows that brains take uncertainty into account when generating decisions. This suggests that neural activity represents beliefs probabilistically. Two proposed classes of probabilistic codes that are broadly compatible with physiological observation include sampling based codes (SBCs) and probabilistic population codes (PPCs). A critical distinction between these codes is the precise manner in which uncertainty is represented. In SBCs, higher uncertainty is associated with greater neural variability, while in PPCs, lower uncertainty is associated with higher firing rates and amplitude encoding of confidence.

A recent study [1] used two different manipulations of spatial information (contrast and stimulus spread) to demonstrate that super Poisson variance was (a) well fit with a divisive normalization model and (b) could be used to decode stimulus uncertainty, i.e. posterior variance. This finding has been used to suggest the presence of SBCs, rather than PPCs (which were not analyzed), for orientation in V1 simple cells. Here we show that this is not the case by deriving a normative model of divisive normalization that maintains a linear PPC in the presence of stimulus spread. This model demonstrates that (1) in the absence of stimulus spread, there is no systematic relationship between excess gain variability and stimulus uncertainty and (2) in the presence of stimulus spread all the information about stimulus uncertainty can be extracted linearly from the population response. In this setting, gain variability provides a redundant and relatively unreliable estimate of stimulus uncertainty even when estimated from many samples.

Finally, we argue that irreducible uncertainty like that modeled by stimulus spread is present in every sensory task since the true value of the stimulus is never observed by the brain. Thus, brains utilizing linear PPC representations should employ divisive normalization ubiquitously and correlation structures should be dominated by patterns associated with gain variability.

3-106. Fast dopamine release signals reward rate to mediate vigor

Ritwik Niyogi[1,2] RITWIK7@GMAIL.COM
Nathaniel Daw[3] NDAW@PRINCETON.EDU

[1]University College London
[2]Max Planck UCL Centre for Computational Psychiatry and Ageing Research
[3]Princeton University

Mesolimbic dopamine release is believed to express motivation by increasing vigor: the propensity to act quicker. But unlike phasic dopaminergic spiking, which is linked to reward prediction error and trial- and-error learning, there is no clear link between dopamine release (which can be distinct from spiking), decision variables, and behavioral vigor. We investigate this question using computational modeling of behavioral vigor and fast-timescale dopamine release dynamics reported from recordings in the nucleus accumbens (Mohebi et al., 2019). Normative analysis (Niv. et al., 2007) implies that the estimated reward rate should govern vigor: animals should act faster when they believe that reward rate is higher. Although this idea has been discussed in terms of putative slow, tonic dopamine responses, we suggest that the more detailed dynamics of reward rate estimation provide a parsimonious formal interpretation for recent observations of faster-timescale peri-event dynamics in nucleus accumbens release, previously argued to track value. This can also be understood as a running reward rate estimate (and hence linked to behavioral vigor), because reward rate can be estimated dynamically by filtering event-related anticipatory signals related to value expectation (i.e., temporal difference prediction errors), rather than averaging rewards actually received. We show that reported peri-event dopamine release to task events superimposes at least two timescales of filtering of such error responses, fast and slow. The slower time constant has the same order of magnitude as that relating behavioral vigor to experienced reward. While this account moves us closer to a unified account of dopamine's dual functions in learning and motivation, a key obstacle remains: that the prediction errors implicit in the dLight release and explicit in phasic spiking do not always coincide. We discuss whether this discrepancy might be explained by local control of release or by nonlinearities in the release function.

3-107. Probabilistic jPCA: a constrained model of neural dynamics

Benjamin Antin[1,2] BENJAMINANTIN1@GMAIL.COM
Krishna Shenoy[3] SHENOY@STANFORD.EDU
Scott Linderman[4] SCOTT.LINDERMAN@GMAIL.COM

[1]Columbia University
[2]Neuroscience
[3]Boston University
[4]Stanford University

One view of nonhuman primate (NHP) and human motor cortex posits that neural activity evolves according to simple dynamical rules Shenoy2013-fc. In many cases, these recordings contain low-dimensional structure, meaning that much of the neural variability can be explained by a small number of latent factors which evolve smoothly over time. To extract these latent factors, Churchland, Cunningham and colleagues proposed jPCA, which projects the data onto planes which capture rotations in state-space. They found that neural trajectories from NHP dorsal premotor (PMd) and primary motor (M1) cortex rotate in state-space during reaches.

Since Churchland, Cunningham and colleagues published jPCA, it has become a standard technique for checking whether rotational dynamics are present in neural recordings. Rotational dynamics have been reported in human motor cortex during reaching and speech, and in artificial networks trained to reproduce muscle activity.

jPCA provides a recipe for extracting rotations, but it does not give a generative model for the data. We cast jPCA as a linear dynamical system (LDS) where the dynamics are constrained to be rotational, which we term "probabilistic jPCA" (pjPCA). We also propose a novel Gibbs sampling algorithm for learning the parameters under the constraint that the dynamics represent rotations. Our model allows several types of analyses not possible with traditional jPCA. First, we obtain uncertainty estimates over the latent factors. Second, our method admits hierarchical variants in which a different dynamics matrix is fit to each trial condition (or possibly each individual trial). Using this hierarchical model, we uncover variability in the frequency of state-space rotations under different trial conditions. Overall, this provides a novel generalization of a well-known technique in systems neuroscience which inherits the benefits typically associated with generative models: model comparison, uncertainty estimation, and hierarchical extensions.

3-108. A normative hippocampus model: robustly encoding variables on smooth manifolds using spiking neurons

Paxon Frady[1,2] PAXON.FRADY@GMAIL.COM
Friedrich Sommer[1] FSOMMER@BERKELEY.EDU

[1]University of California, Berkeley
[2]Intel

Different brain regions extract and manipulate behavioral variables, which are often continuous and organized on smooth manifolds (Stringer et al., 2019), such as location in an environment. A useful encoding scheme for these variables has to fulfill two fundamental requirements: (1) Distinguishability: Points far apart on a manifold, should be encoded or "indexed" by dissimilar (independent or orthogonal) patterns so that the probability of confusion is low. (2)Smoothness: Points close to each other should be encoded by similar patterns such that the code is noise tolerant and points can be represented by interpolation of neighbors. Here we propose a normative encoding model with properties (1)-(2), leveraging timing of individual spikes. Our approach is based on complex-valued connectionist models for symbolic reasoning (Plate, 2003),and recent theoretical results on how complex state spaces can be mapped onto the dynamics of spiking neural networks through a phase-to-timing mapping (Frady and Sommer, 2019). The resulting model "tiles" a manifold such that each location is represented by a unique neural activity vector that can be used as an index (1), but also in a way that transitions between neighboring tiles are smooth (2). This is accomplished naturally in the complex domain, with the sinc function as a universal similarity kernel. Our approach is demonstrated in a model of hippocampus neurons encoding variables such as the location of the animal and also head direction, satisfying conditions(1) and (2). Interestingly, prominent coding properties in hippocampus, such as place cells and spike-phase precession are naturally reproduced by our normative model without building it in a-priori. Further, our model helps to explain recent observations of conjunctional receptive fields(Høydal et al., 2019), and how to relate computation with manifold structure observed in neural recordings (Low et al., 2018). Specifically, the model makes concrete predictions how the intrinsic structure of neural recordings is influenced by: (i) manifold dimensionality (number of variables),(ii) manifold smoothness and resolution (number of tiles), and (iii) redundant neural dimensions for robustness.

3-109. Optimality theories and statistical analysis of neural systems

Wiktor Mlynarski[1] WMLYNARS@IST.AC.AT
Michal Hledik[1] MICHAL.HLEDIK@IST.AC.AT
Thomas Sokolowski[2] SOKOLOWSKI@FIAS.UNI-FRANKFURT.DE
Gasper Tkacik[1] GASPER.TKACIK@IST.AC.AT

[1]Institute of Science and Technology Austria
[2]Frankfurt Institute for Advanced Studies

Normative theories and statistical inference provide complementary approaches for the study of neural coding and behavior. A normative theory postulates that organisms have adapted to efficiently solve essential tasks, and proceeds to mathematically work out testable consequences of such optimality; parameters that maximize the hypothesized function can be derived ab initio, without reference to experimental data. In contrast, statistical inference focuses on efficient utilization of data to learn model parameters, without reference to any a priori notion of biological function, utility, or fitness. Traditionally, these two approaches were developed independently and applied separately.

In this work we provide a unifying account of statistical inference and normative theories. We develop a Bayesian framework that embeds a normative theory into a family of maximum-entropy "optimization priors." This family defines a smooth interpolation between a data-rich inference regime (characteristic of "bottom-up" statistical models), and a data-limited ab inito prediction regime (characteristic of "top-down" normative theory). These optimization priors establish a connection between theories of biological optimality and the rich toolbox of statistical data analysis.

We introduce our approach using a simple toy model of a neuron optimized to transmit information about sensory stimuli. We further demonstrate its generality by analyzing receptive fields in the visual cortex, temporal neural filters in the retina, and neural connectivity in C. elegans. These examples illustrate that our framework allows to provide statistically rigorous answers to questions about biological function and optimality. We argue that the flexibility it affords is essential to address a number of fundamental challenges relating to inference and prediction in complex, high-dimensional biological problems.

3-110. Two distinct mechanisms for time perception are selected by context

Sebastian Reinartz REINARTZ.SEBASTIAN@GMAIL.COM
Maria Ravera MRAVERA@SISSA.IT
Mario Fontanini MARIO.FONTANINI1992@GMAIL.COM
Mathew E Diamond DIAMOND@SISSA.IT

International School for Advanced Studies (ISAS/SISSA)

Every day we judge the duration of countless sensory events, implicitly or explicitly. Time perception has an intimate connection to the sensory features of the event – when rats judge the duration of a $<$1-second vibration applied to the whiskers, increasing vibration intensity leads to increasing perceived duration (https://doi.org/10.1101/2020.08.02.232801). When spike trains from vibrissal somatosensory cortex (vS1) are modelled as input to a leaky integrator, resulting neurometric curves match observed psychometric curves. Furthermore, optogenetic manipulation of vS1 alters perceived duration (Authors, Cosyne 2019). However, when sensory events are marked by clear onset/offset signals, does the same integrative mechanism hold? We trained rats to classify stimuli, defined by 7 possible durations and 5 intensities, as 'long' or 'short.' Each stimulus was either (i) "uniform," a continuous, noisy vibration identical to those of earlier studies, or (ii) "flanked," a noisy vibration where first and final 25ms were amplified, yielding clear onset/offset. In uniform sessions, perceived duration was biased by stimulus intensity and by optogenetic excitation of vS1, consistent with integration of vS1 sensory drive. In flanked sessions, the bias evoked by intensity and optogenetic excitation of vS1 disappeared. Individual stimuli were processed according to their context; in flanked context (85% of trials flanked, remaining 15% uniform), uniform stimuli were judged as if flanked – no bias was imposed by vibration intensity. Symmetrically, in uniform-context sessions, flanked stimuli (the 15% of trials) showed an intensity bias. This suggests a top-down selection of the processing mechanism, as opposed to a bottom-up selection wherein processing mechanisms would vary trial-to-trial according to the features of the individual stimulus. The findings suggest gating between two distinct mechanisms for duration perception: 1) integration of vS1 sensory drive for uniform stimuli, and 2) integration of some input distinct from the sensory drive for stimuli with highly salient onset/offset.

3-111. Optimal sparseness depends on task complexity, input dimensionality and training set size

Ramon Nogueira RN2446@COLUMBIA.EDU
Stefano Fusi SF2237@COLUMBIA.EDU

Columbia University

A key issue in systems and computational neuroscience is to identify the features of neural representations that can generate flexible and robust behavior. High dimensional representations allow linear readouts to perform a large number of different complex discrimination tasks while low dimensional representations are advantageous for generalization to unseen experimental conditions. Previous theoretical studies have shown that the sparseness of neural representations can efficiently control this generalization-discrimination trade off. Here we characterize the role of sparseness on the performance of complex and simple discrimination tasks in more realistic networks that contain a layer of randomly connected neurons. We extended previous results by considering continuous ReLu units with different input structures, number of training samples, noise levels and neural architectures. Interestingly, in the case of complex tasks, for a low number of training samples, 10% - 20% of active neurons (coding level) produced the highest task performance. When several training samples were used, coding levels of up to ~50% were equally optimal. This observation could explain the difference between the low coding level observed in the hippocampus, where the short memory limits the number of samples used for training, and the relatively high coding level of the cortex, where slower learning might allow to consider a larger number of samples. Moreover, we observed that for high noise levels and many training examples, reading out directly from the input layer produced a higher performance than reading out from the high dimensional intermediate layer, indicating that high-dimensionality can be harmful for generalization when the noise is high. Our work provides a novel theoretical framework that can predict the optimal sparseness level for different types of tasks, number of training samples, input dimensionality and neural architectures. This framework could be used to explain the variety of coding levels found across cortical layers and brain regions.

3-112. Functional segregation of the marmoset auditory cortex by data-driven decomposition of responses to naturalistic sounds with wide-field optical imaging

Yueqi Guo[1,2] GUOYUEQI912@GMAIL.COM
Xindong Song[1] SONGXINDONG@JHMI.EDU
Chenggang Chen[1] CHENG-GANG.CHEN@JHU.EDU
Xiaoqin Wang[1] XIAOQIN.WANG@JHU.EDU

[1]Johns Hopkins University
[2]Biomedical Engineering

Primates have developed a signature structure in the auditory cortex that is generally characterized by 3 layers of hierarchies – core, belt and parabelt organized along the medial-lateral axis. While previous studies have extensively evaluated neural activities in the core region of auditory cortex driven by pure tone stimuli, how the entire auditory cortex is spatially organized to process naturalistic sounds with rich acoustic features, especially in the secondary auditory areas (belt and parabelt), is still poorly understood. In this study, we sought to investigate this fundamental question in the common marmoset (Callithrix jacchus), a highly vocal New World monkey species that has garnered considerable interest in recent years as a promising non-human primate model in neuroscience. We developed chronic optical imaging techniques to obtain wide-field intrinsic and calcium signals from the auditory cortex in awake marmosets. Our results showed that naturalistic sounds elicited responses in both primary and secondary auditory regions, and that the specific sound can be decoded from wide-field calcium signal with a high accuracy. With an unsupervised matrix decomposition method [Norman-Haignere et al., 2015], we identified canonical functional components (subregions) in the auditory cortex that are sensitive to different acoustic features, revealed by correlation analysis between component response profiles and stimulus features. Interestingly, we found that the anterior core region that is sensitive to low frequencies (RT) responds to music and speech more strongly than other categories, which could potentially be a predecessor for the music or speech area in humans. Our results revealed a functional organization in the auditory cortex of primates that is beyond tonotopy, and provided insights for understanding how the information is processed hierarchically.

3-113. World-centered and self-centered spatial representation during head orienting movements in the rat frontal orienting field

Liujunli Li[1,2] LL3462@NYU.EDU
Jingjie Li[1] JINGJIE.LI@NYU.EDU
Ce Ma[1] CM4913@NYU.EDU
Jeffrey Erlich[1] JERLICH@NYU.EDU

[1]New York University Shanghai
[2]Center for Neural Science

Navigation in a complex world requires integration of world-centered and self-centered spatial representation. Self-centered direction tuned neurons have been extensively characterized in the mammalian frontal cortex, including the primate frontal eye fields (FEF) and the rodent frontal orienting fields (FOF). However, whether world-centered representation is also encoded in frontal cortex, and how the two types of information are integrated remain unclear. To simultaneously study world- and self-centered representation in freely-moving animals beyond 2-alternative forced choices, we developed a novel spatial orienting apparatus. On our newly designed port wall animals can move between 7 ports in 6 possible directions, which allows us to compare the same movements from different start positions or different movements to the same target. Here, we recorded electrophysiologically from the FOF of rats performing delayed visually-guided head-orienting. On each trial, rats were first cued to fixate in a central start port, and then guided visually to poke into one of the 6 target ports. Consistent with past recordings in the primate FEF and rat FOF, we found self-centered direction-tuned neurons, which fired before and during the movement. Surprisingly, we also found world-centered position-tuned neurons. Similar to hippocampal place cells, many of these position neurons encoded the current position, preferring a specific port as both the start and target port, irrespective of proximal movement direction. On the population level, direction tuning for the target port emerged earlier than position tuning, suggesting that motor plans are preferentially encoded in self-centered rather than world-centered reference frames. Together, our results suggest that the FOF may play a role in integrating multiple , rather than simply movement planning.

3-114. The emergence and function of cortical offset responses in sound termination detection

Magdalena Solyga[1,2] MAGDALENA.SOLYGA@UNIBAS.CH
Tania Barkat[1] TANIA.BARKAT@UNIBAS.CH

[1]University of Basel
[2]Department of Biomedicine

Offset responses in auditory processing appear after a sound terminates. In the peripheral auditory system, neuronal circuits within the cochlear and superior paraolivary nuclei play a central role in generating offset responses. In the mouse central auditory system, a recent study showed that neurons in the anterior auditory field (AAF) have significantly stronger responses to sound termination than those in the primary auditory cortex. The behavioural relevance of cortical offset responses, as well as the mechanisms that generate and drive them within AAF, remain unknown. To understand the role of AAF offset responses in sound perception, we asked whether detecting the end of a sound would be influenced by them. Using optogenetics and behavioural paradigms, we first showed that preventing AAF offset responses decreases the mouse performance to detect sound termination. We also demonstrated that bigger AAF offset responses facilitate the detection of when a sound ends. We then asked how cortical offset responses occur. By combining electrophysiological recordings in the cortex and thalamus with antidromic stimulation, we found that 80% of AAF inputs come from onset-offset responsive thalamic cells. However, our results also demonstrate that cortical offset responses are amplified in comparison to the thalamic ones. For example, offset response amplitudes are significantly increasing with sound duration in the AAF but not in the MGB. Also, offset responses to white noise stimulation are present in the AAF but not in the MGB. Together, this indicates that cortical offset responses are partially inherited from the preceding nucleus along the auditory pathway, but that they can also be generated de novo within cortical circuits. Overall, our results reveal the importance of cortical offset responses in coding and perceiving sound termination. They might play a crucial role for advanced processing such as detecting changes within temporally discontinuous sounds like speech and vocalization.

3-115. A cortical-hypothalamic circuit decodes social rank and promotes dominance behavior

Nancy Padilla Coreano[1] NANCYPADILLA.UPR@GMAIL.COM
Kanha Batra[2,1] KBATRA@ENG.UCSD.EDU
Makenzie Patarino[1] MPATARINO@SALK.EDU
Zexin Chen[3] CZXSJTU@GMAIL.COM
Rachel R. Rock[1,4] RACHEL.ROCK@UCSF.EDU
Ruihan Zhang[4] RUIHANZ@MIT.EDU
Sebastien Hausmann[1,5] SEBASTIEN.HAUSMANN@EPFL.CH
Javier C. Weddington[4] JAVIERCW@STANFORD.EDU
Reesha Patel[1] RPATEL@SALK.EDU
Yu Eva Zhang[2] YUZ098@UCSD.EDU
Hao-Shu Fang[3] FHAOSHU@GMAIL.COM
Laurel R. Keyes[1] LKEYES@SALK.EDU
Avraham Libster[1] AVI.LIBSTER@GMAIL.COM
Gillian Matthews[1] GILLIANMATTHEWS85@GMAIL.COM
James P. Curley[6] CURLEY@UTEXAS.EDU
Ila R. Fiete[4] FIETE@MIT.EDU
Cewu Lu[3] LUCEWU@SJTU.EDU.CN
Kay M. Tye[1] TYE@SALK.EDU

[1]Salk Institute for Biological Studies
[2]University of California, San Diego
[3]Shanghai Jiao Tong University
[4]Massachusetts Institute of Technology
[5]Ecole polytechnique federale de Lausanne
[6]University of Texas at Austin

How do we know our social rank? Most social species, from insects to humans, self-organize into social dominance hierarchies. The establishment of social ranks serves to decrease aggression, conserve energy, and maximize survival for the entire group. Despite dominance behaviors being critical for successful interactions and ultimately, survival, we have only begun to learn how the brain represents social rank and guides behavior based on this representation. The medial prefrontal cortex (mPFC) has been implicated in the expression of social dominance in rodents and in social rank learning in humans. Yet precisely how the mPFC encodes rank and which

circuits mediate this computation is not known. We developed a trial-based social competition assay in which mice compete for rewards, as well as a computer vision tool to track multiple, unmarked animals. With the development of a deep learning computer vision tool (AlphaTracker) and wireless electrophysiology recording devices, we have established a novel platform to facilitate quantitative examination of how the brain gives rise to social behaviors. We describe nine behavioral states during social competition that were accurately decoded from mPFC ensemble activity using a hidden Markov model combined with generalized linear models (HMM-GLM; Escola et al., 2011; Calhoun et al., 2019). Population dynamics in the mPFC were predictive of social rank and competitive success. This population-level rank representation translated into differences in the individual cell responses to task-relevant events across ranks. Finally, we demonstrate that mPFC cells that project to the lateral hypothalamus contribute to the prediction of social rank and promote dominance behavior during the reward competition. Thus, we reveal a cortico-hypothalamic circuit by which mPFC exerts top-down modulation of social dominance.

3-116. Embedding disentangled dynamics from behavioral videos

Changhao Shi[1] CSHI@UCSD.EDU
Shahar Levy[2] SHAHAR.LEVY86@GMAIL.COM
Maria Lavzin[2] MARIA.LAVZIN@GMAIL.COM
Jackie Schiller[2] JACKIE@TECHNION.AC.IL
Gal Mishne[1,3] GMISHNE@UCSD.EDU

[1]University of California, San Diego
[2]Technion - Israel Institute of Technology
[3]HDSI

Automated analysis of behavioral videos is a subject of increasing interest in the neuroscience community, in the goal to understand the relationship between neural activity and behavior. Leveraging recent advances in deep learning, we develop a novel end-to-end deep autoencoder framework to learn an informative low-dimensional embedding of high-dimensional video frames. Compared with traditional analysis methods or recent deep learning methods (e.g. DeepLabCut), our method encodes all visible behavioral features and generate both fine-grained embedding and coarse event labeling. In order to learn a consistent embedding space of the behavioral poses across different animals and different visual layouts, the network incorporates a disentangling mechanism that separates the behavioral dynamics (pose) from non-behavioral variables. In addition to the fine-resolution embedding of the pose component, the network learns behavioral states via a latent Gaussian mixture model (GMM) where each cluster corresponds to a behavioral state. Furthermore, the network can be trained independently on behavioral videos in an unsupervised manner or in a semi-supervised manner using a small number of labeled frames. We demonstrate the efficacy of our approach on a large dataset of motor task learning in mice with front and profile video views. The head-fixed mice are trained to reach for a food pellet following an auditory cue. This complex behavioral task mimics distal hand movements in primates and involves fine movements of the forepaw and digits. Our method learns an informative embedding of dynamics across animals and learning.

3-117. Modelling functional wiring and processing from retinal bipolar to ganglion cells

Klindt David[1] KLINDT.DAVID@GMAIL.COM
Philipp Berens[2] PHILIPP.BERENS@UNI-TUEBINGEN.DE
Thomas Euler[3] THOMAS.EULER@CIN.UNI-TUEBINGEN.DE
Cornelius Schroder[1] CORNELIUS.SCHROEDER@GMAIL.COM
Anna Vlasits[1] ANNA.VLASITS@GMAIL.COM
Katrin Franke[1] KATRIN.FRANKE@TUEBINGEN.MPG.DE

[1]University of Tubingen
[2]Institute for Ophthalmic Research, University of Tubingen
[3]University of Tuebingen

Although the processing of visual information in the retina has been studied in detail, the underlying functional connectivity is not yet completely understood. While many specific circuits are well-characterized (e.g. the rod photoreceptor pathway), a comprehensive picture of how these microcircuits work together to form the retinal network is still lacking. Furthermore, connectomic information, which could help dissect the functional underpinnings of the retina, is not yet fully leveraged. The integration of different datasets and data sources to (computational) models is a key challenge to elucidate the processing of visual information in the retina. A step towards a comprehensive understanding of the retinal network was made in a recent publication [1], which suggests a biophysically-constrained bipolar and amacrine cell network model (BCN model) of light processing in the mouse

inner retina, enabling in silico experiments. Here, we extended this model to predict the responses of previously characterized mouse retinal ganglion cell (RGC) types to full-field light stimulation [2]. Specifically, we tested how bipolar cell glutamate release can be combined in an additional linear nonlinear model to predict RGC output (BCN-LN). We show that recordings of full-field stimulation combined with mechanistically detailed modelling allowed us to predict connectivity between cell types, as well as to investigate the role of inhibitory feedback and feedforward AC modulation. In summary, this work shows how a machine-learning approach informed by biological structures can produce interpretable and accurate predictions about neural connectivity and circuit functions, thus explaining the hierarchical emergence of functional RGC diversity.

3-118. Adversarial training aligns invariances between artificial neural networks and biological sensory systems

Jenelle Feather[1,2] JFEATHER@MIT.EDU
Alex Durango[3] DURANGOA@MIT.EDU
Guillaume Leclerc[1] LECLERC@MIT.EDU
Aleksander Mądry[1] MADRY@MIT.EDU
Josh McDermott[1] JHM@MIT.EDU

[1]Massachusetts Institute of Technology
[2]Brain and Cognitive Sciences
[3]Stanford University

Progress in deep learning has produced computational models that can perform complex tasks as well as biological organisms and predict auditory and visual brain responses better than previous models. The transformations learned by these models are thought to capture invariances like those of biological perceptual systems. However, several strands of recent work have highlighted discrepancies between artificial neural networks and humans that suggest differences between their invariances. One striking divergence is evident in “model metamers” – stimuli that produce nearly identical activations within an artificial neural network, but that physically differ along dimensions to which the model is invariant. Model metamers generated to match deep activations of natural images or audio are generally unrecognizable to humans despite producing the same class predictions as the natural signal in the network used to generate them, indicating that the model is invariant to stimulus dimensions to which humans are not. Another divergence is evident with “adversarial examples” – stimuli that are perturbed by an amount that is imperceptible to humans but that produce a different class label by the network. To explore how these two phenomena are related, we measured the recognizability of model metamers generated from networks trained to be robust to adversarial perturbations. In both the visual and auditory domain, robustly trained networks produced model metamers that were more recognizable to humans than those generated from standard-trained models. These differences between models were not evident using the more common model evaluation metric of brain response predictions from network activations. The results suggest that making models more invariant to perturbations that are imperceptible to humans (reducing adversarial examples) reduces other model invariances that humans lack, better aligning deep neural networks to biological sensory systems. The results highlight the value of model metamers in revealing differences between models that are not apparent with other metrics.

3-119. Free recall scaling laws and short-term memory effects in latching attractor networks

Vezha Boboeva[1,2] VBOBOEVA@IC.AC.UK
Alberto Pezzotta[3] ALBERTO.PEZZOTTA@CRICK.AC.UK
Claudia Clopath[1] CCLOPATH@IC.AC.UK

[1]Imperial College London
[2]Department of Bioengineering
[3]The Francis Crick Institute

Despite the complexity of human episodic memory, paradigms such as free recall have revealed surprisingly robust qualitative and quantitative characteristics, such as laws governing recall capacity. Although abstract random matrix models could explain such scaling laws, the possibility of their implementation in large networks of interacting units has so far remained unexplored. We study an attractor network model of long term memory endowed with firing rate adaptation and global inhibition, and find that under appropriate conditions, the latching ability of the network, i.e. its transitioning behaviour from memory to memory, is constrained by limit cycles that prevent the network from recalling all memories, with scaling similar to what has been found in experiments. In addition, when the model is supplemented with additional ingredients, such as a hetero-associative learning rule, complementing the standard auto-associative learning rule, as well as short term synaptic facilitation, our model reproduces other

findings in the free recall literature such as serial position effects, contiguity and forward asymmetry effects in the probability to transition from memory to memory, as well as the semantic effects found to guide memory recall. The model is consistent with a broad series of manipulations aimed at gaining a better understanding of the variables that affect recall, such as the role of rehearsal, presentation rates, and (continuous/end-of-list) distractor conditions.

3-120. Hebbian plasticity in parallel synaptic pathways: A circuit mechanism for systems memory consolidation

Denis Alevi[1] DENIS.ALEVI@BCCN-BERLIN.DE
Michiel Remme[2] MICHIEL.REMME@GMAIL.COM
Urs Bergmann[2] URSBERGMANN@GMAIL.COM
Susanne Schreiber[2] S.SCHREIBER@HU-BERLIN.DE
Richard Kempter[2,3] R.KEMPTER@BIOLOGIE.HU-BERLIN.DE
Henning Sprekeler[1] H.SPREKELER@TU-BERLIN.DE

[1]Technische Universitat Berlin
[2]Humboldt-Universitat zu Berlin
[3]Department of Biology - Theoretical Neuroscience

Systems memory consolidation can transfer declarative memories which initially depend on the hippocampal formation into long-term memory traces in neocortical networks. This consolidation process is thought to involve the replay of recently acquired memories, but the cellular and network mechanisms that mediate the memory transfer are poorly understood.

Here, we suggest that systems memory consolidation could arise from Hebbian plasticity in networks with parallel synaptic pathways — two ubiquitous features of neural circuits in the brain. We explore this hypothesis in a computational model and show that it is in quantitative agreement with lesion studies in rodents and that the mechanism can be hierarchically iterated to yield power-law forgetting as observed in psychophysical studies in humans.

Further, we show that the proposed mechanism enables the extraction of semantic generalizations from individual memories and forgetting of episodic detail during consolidation. We find that while neural replay of specific experiences is necessary for the extraction of semantic generalizations, random neural activity is sufficient for the consolidation of already generalized memory content.

In summary, our theory may provide an inrode to the mechanistic understanding of memory reorganization and transformation by bridging spatial scales from single cells to cortical areas and time scales from milliseconds to years.

Author Index